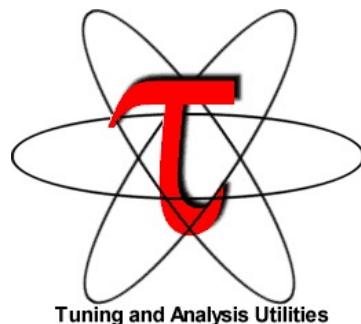


# TAU

August 11, 2021, 3:40pm – 4:15pm CT  
ATPESC workshop via Zoom

Sameer Shende  
Performance Research Laboratory, OACISS, University of Oregon  
[http://tau.uoregon.edu/tau\\_atpesc21.pdf](http://tau.uoregon.edu/tau_atpesc21.pdf)



# TAU: Quickstart Guide

## Setup:

- % module load tau [ on Theta, use aprun instead of mpirun, and on ThetaGPU please use:]
- % module use /soft/perf-tools/tau/modulefiles/thetagpu; module load tau

## Profiling with an un-instrumented application:

- MPI: % mpirun -np 64 tau\_exec -ebs ./a.out
  - MPI+OpenMP with Intel 19+: % export TAU\_OMPT\_SUPPORT\_LEVEL=full;  
% mpirun -np 64 tau\_exec -T ompt,v5 -ompt ./a.out
  - Pthread: % mpirun -np 64 tau\_exec -T mpi,pthread -ebs ./a.out
  - Python+MPI+Sampling: % mpirun -np 64 tau\_python -ebs ./a.py
  - Python+MPI+OpenCL: % mpirun -np 64 tau\_python -opencl ./a.py
  - DPC++/SYCL (no MPI): % tau\_exec -T level\_zero,serial -10 ./a.out
  - CUDA+MPI % mpirun -np 64 tau\_exec -T cupti -ebs -cupti ./a.out
- Analysis: % pprof -a -m | more; % paraprof (GUI)

## Tracing:

- Vampir: MPI: % export TAU\_TRACE=1; export TAU\_TRACE\_FORMAT=otf2  
% mpirun -np 64 tau\_exec ./a.out; vampir traces.otf2 &
- Chrome: % export TAU\_TRACE=1; mpirun -np 64 tau\_exec ./a.out; tau\_treemerge.pl;  
% tau\_trace2json tau.trc tau.edf -chrome -ignoreatomic -o app.json  
Chrome browser: chrome://tracing (Load -> app.json)
- Jumpshot: % export TAU\_TRACE=1; mpirun -np 64 tau\_exec ./a.out; tau\_treemerge.pl;  
% tau2slog2 tau.trc tau.edf -o app.slog2; jumpshot app.slog2 &

# Challenges

- With growing hardware complexity, it is getting harder to accurately measure and optimize the performance of our HPC and AI/ML workloads.
- As our software gets more complex, it is getting harder to install tools and libraries correctly in an integrated and interoperable software stack.

# Motivation: Improving Productivity

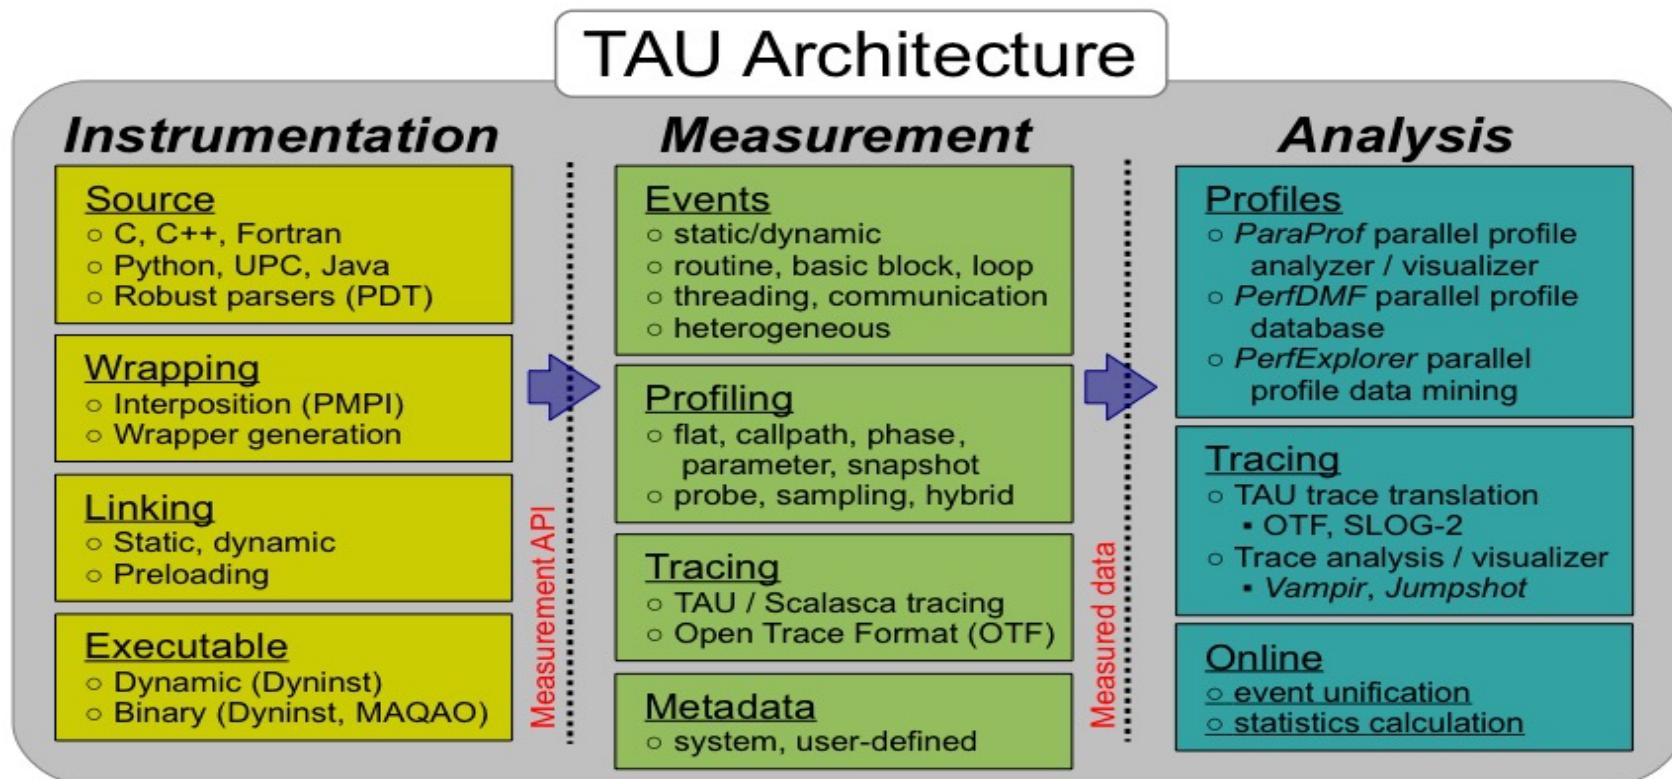
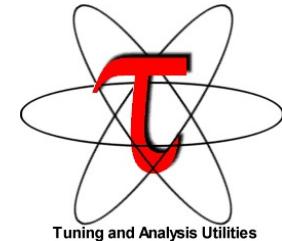
- TAU Performance System<sup>®</sup>:
  - Deliver a scalable, portable, performance evaluation toolkit for HPC and AI/ML workloads
- Extreme-scale Scientific Software Stack (E4S):
  - Delivering a modular, interoperable, and deployable software stack
  - Deliver expanded and vertically integrated software stacks to achieve full potential of extreme-scale computing
  - Lower barrier to using software technology (ST) products from ECP
  - Enable uniform APIs where possible

# TAU Performance System®

## Parallel performance framework and toolkit

Supports all HPC platforms, compilers, runtime system

Provides portable instrumentation, measurement, analysis



# TAU Performance System

## Instrumentation

- Fortran, C++, C, UPC, Java, Python, Chapel, Spark
- Automatic instrumentation

## Measurement and analysis support

- MPI, OpenSHMEM, ARMCI, PGAS, DMAPP
- pthreads, OpenMP, OMPT interface, hybrid, other thread models
- GPU: Intel oneAPI DPC++/SYCL, AMD ROCm, CUDA, OpenCL, OpenACC, Kokkos
- Parallel profiling and tracing

## Analysis

- Parallel profile analysis (ParaProf), data mining (PerfExplorer)
- Performance database technology (TAUdb)
- 3D profile browser

# Application Performance Engineering using TAU

- How much time is spent in each application routine and outer *loops*? Within loops, what is the contribution of each *statement*? What is the time spent in OpenMP loops? In kernels on GPUs. How long did it take to transfer data between host and device (GPU)?
- How many instructions are executed in these code regions?  
Floating point, Level 1 and 2 *data cache misses*, hits, branches taken? What is the extent of vectorization for loops?
- What is the memory usage of the code? When and where is memory allocated/de-allocated? Are there any memory leaks? What is the memory footprint of the application? What is the memory high water mark?
- How much energy does the application use in Joules? What is the peak power usage?
- What are the I/O characteristics of the code? What is the peak read and write *bandwidth* of individual calls, total volume?
- How does the application *scale*? What is the efficiency, runtime breakdown of performance across different core counts?

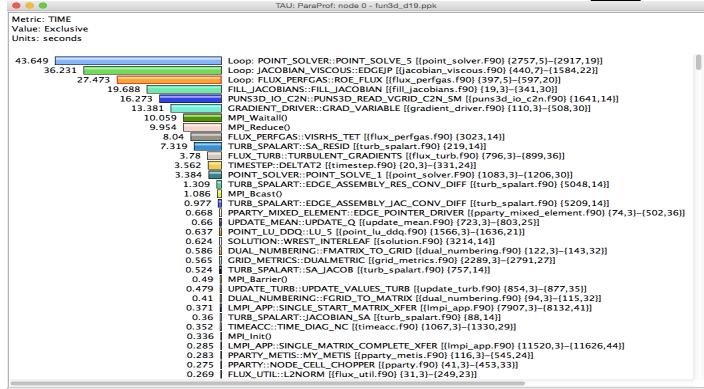
# Instrumentation

## Add hooks in the code to perform measurements

- **Source instrumentation using a preprocessor**
  - Add timer start/stop calls in a copy of the source code.
  - Use Program Database Toolkit (PDT) for parsing source code.
  - Requires recompiling the code using TAU shell scripts (tau\_cc.sh, tau\_f90.sh)
  - Selective instrumentation (filter file) can reduce runtime overhead and narrow instrumentation focus.
- **Compiler-based instrumentation**
  - Use system compiler to add a special flag to insert hooks at routine entry/exit.
  - Requires recompiling using TAU compiler scripts (tau\_cc.sh, tau\_f90.sh...)
- **Runtime preloading of TAU's Dynamic Shared Object (DSO)**
  - No need to recompile code! Use `aprun tau_exec ./app` with options.

# Profiling and Tracing

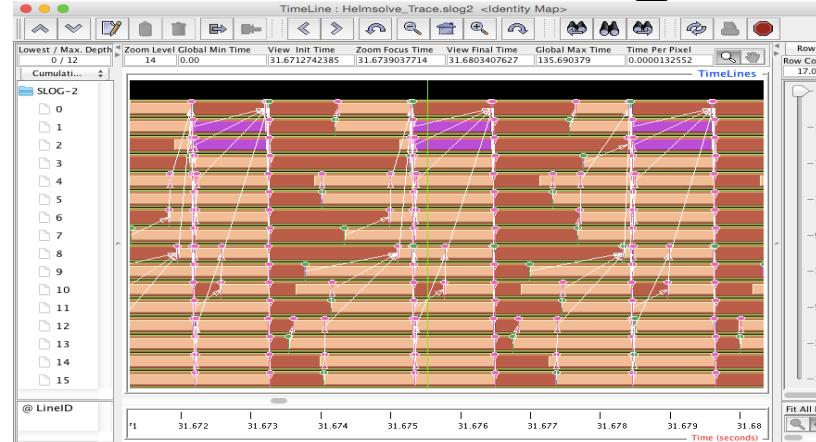
## Profiling



- Profiling shows you **how much** (total) time was spent in each routine
- Profiling and tracing

**Profiling shows you how much (total) time was spent in each routine**  
**Tracing shows you when the events take place on a timeline**

## Tracing



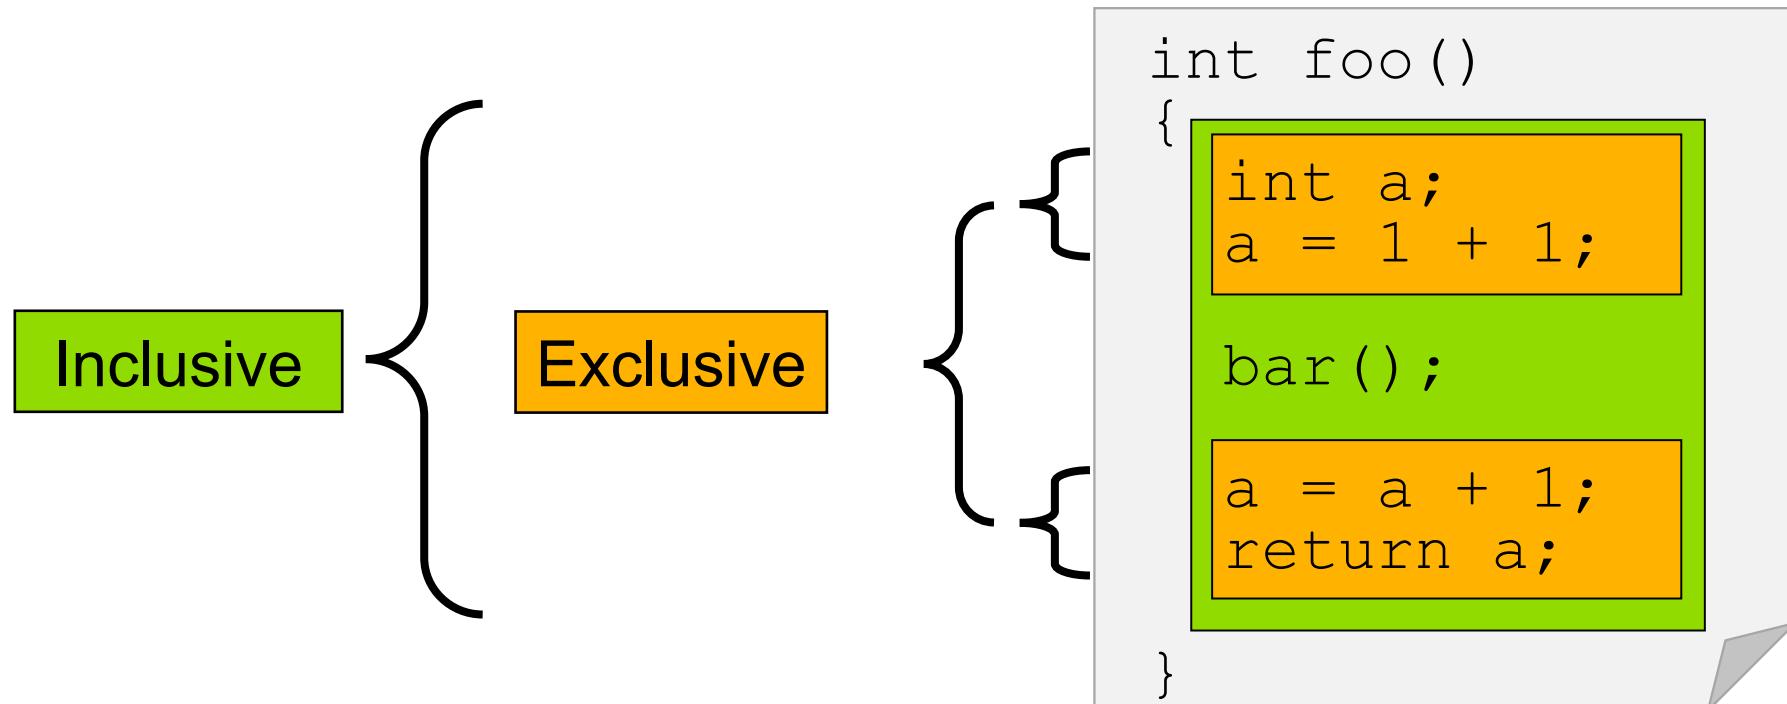
- Tracing shows you when the events take place on a timeline

# Instrumentation

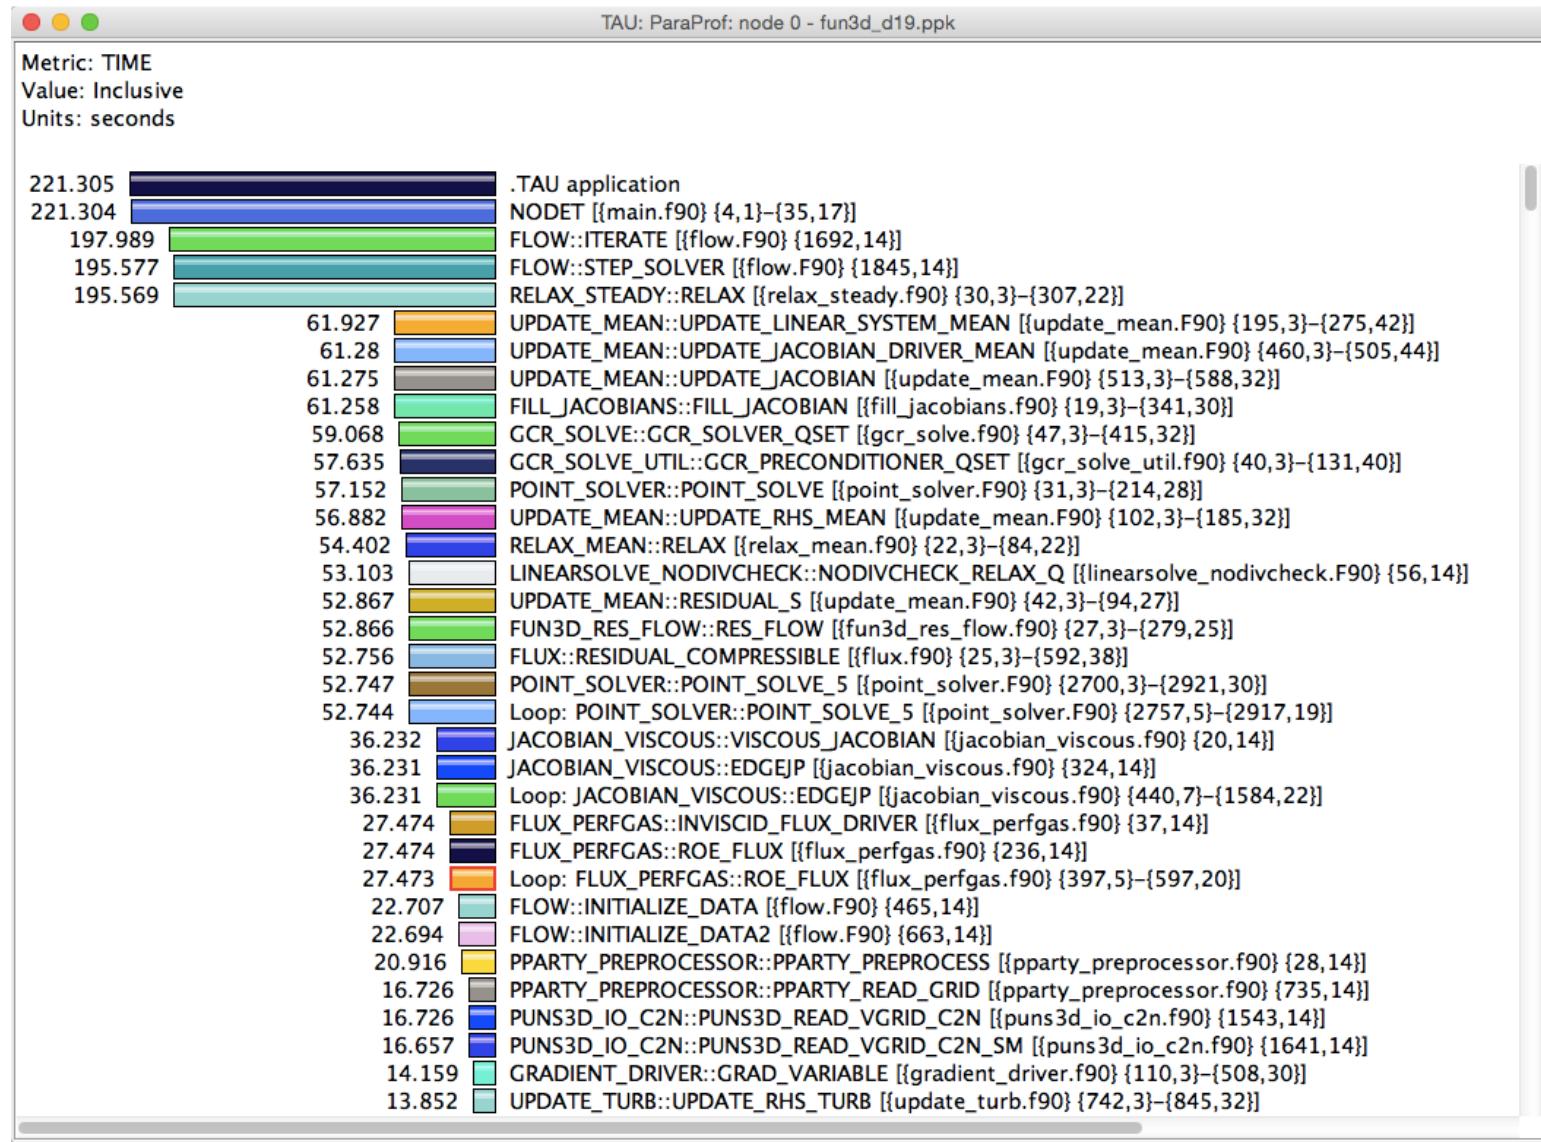
- Direct and indirect performance observation
- Instrumentation invokes performance measurement
- Direct measurement with *probes*
- Indirect measurement with periodic sampling or hardware performance counter overflow interrupts
- Events measure performance data, metadata, context, etc.
- User-defined events
  - **Interval** (start/stop) events to measure exclusive & inclusive duration
  - **Atomic events** take measurements at a single point
    - Measures total, samples, min/max/mean/std. deviation statistics
  - **Context events** are atomic events with executing context
    - Measures above statistics for a given calling path

# Inclusive vs. Exclusive values

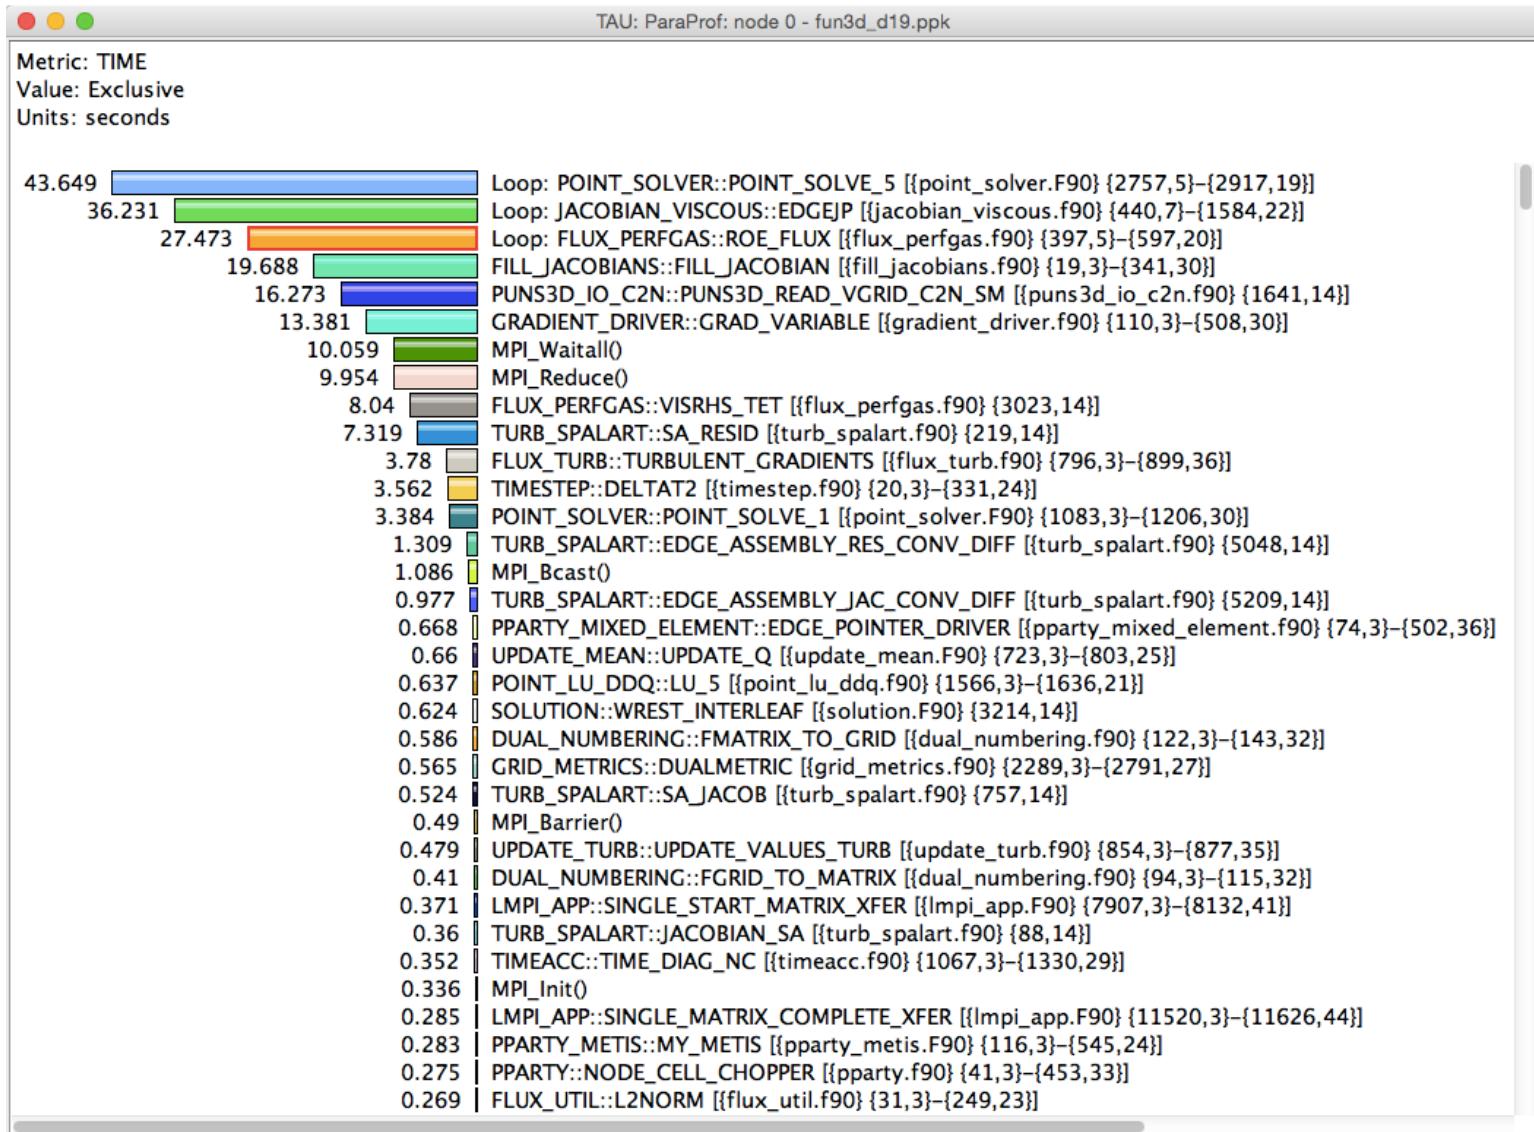
- Inclusive
  - Information of all sub-elements aggregated into single value
- Exclusive
  - Information cannot be subdivided further



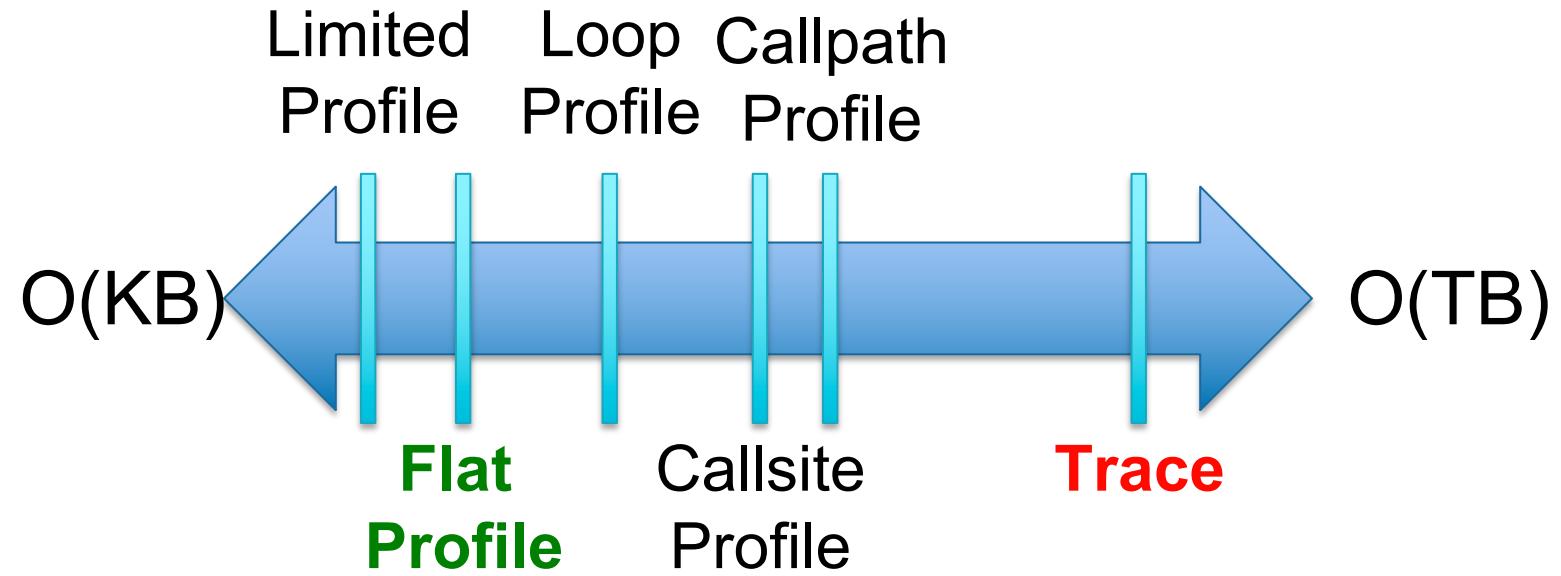
# Inclusive Measurements



# Exclusive Time



# How much data do you want?



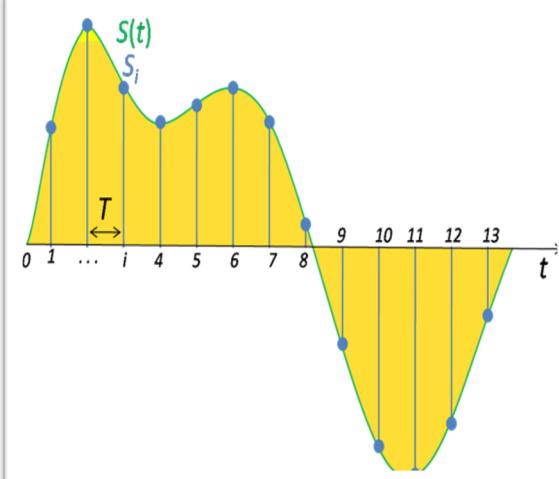
# Performance Data Measurement

## Direct via Probes

```
Call START('potential')
// code
Call STOP('potential')
```

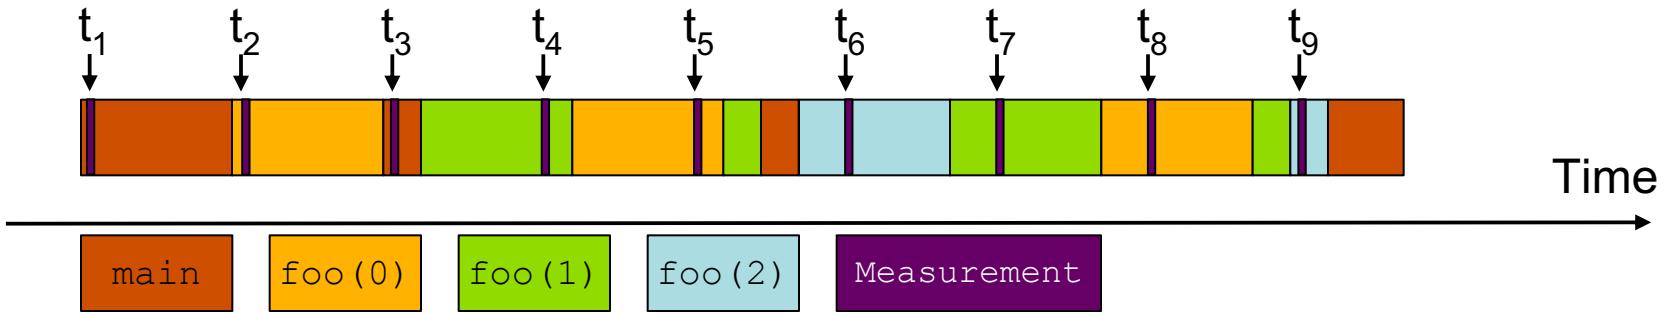
- Exact measurement
- Fine-grain control
- Calls inserted into code

## Indirect via Sampling



- No code modification
- Minimal effort
- Relies on debug symbols (**-g**)

# Sampling



- Running program is periodically interrupted to take measurement
  - Timer interrupt, OS signal, or HWC overflow
  - Service routine examines return-address stack
  - Addresses are mapped to routines using symbol table information
- Statistical inference of program behavior
  - Not very detailed information on highly volatile metrics
  - Requires long-running applications
- Works with unmodified executables

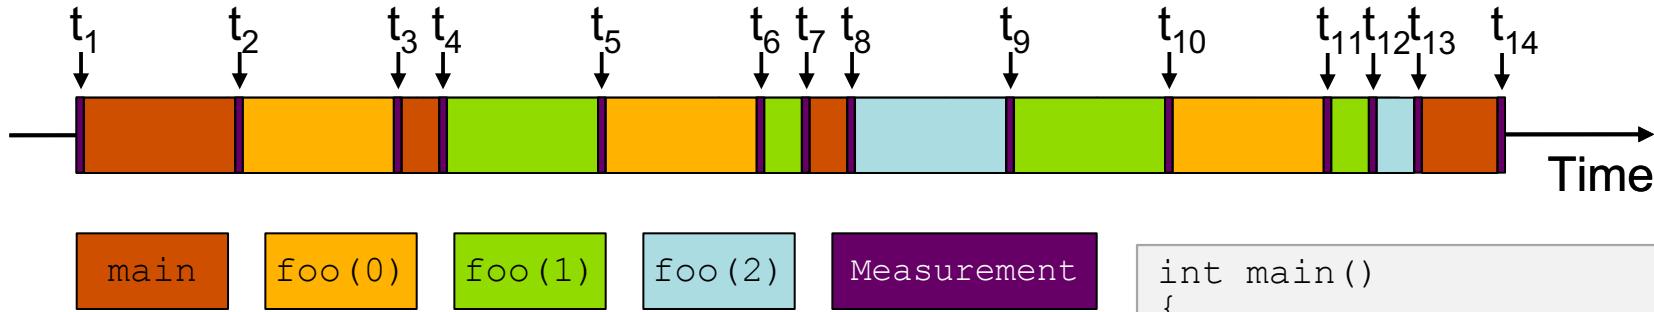
```
int main()
{
    int i;

    for (i=0; i < 3; i++)
        foo(i);

    return 0;
}

void foo(int i)
{
    if (i > 0)
        foo(i - 1);
}
```

# Instrumentation



- Measurement code is inserted such that every event of interest is captured directly
  - Can be done in various ways
- Advantage:
  - Much more detailed information
- Disadvantage:
  - Processing of source-code / executable necessary
  - Large relative overheads for small functions

```
int main()
{
    int i;
    Start("main");
    for (i=0; i < 3; i++)
        foo(i);
    Stop("main");
    return 0;
}

void foo(int i)
{
    Start("foo");
    if (i > 0)
        foo(i - 1);
    Stop("foo");
}
```

# Using TAU's Runtime Preloading Tool: tau\_exec

Preload a wrapper that intercepts the runtime system call and substitutes with another

**MPI**

**OpenMP**

**POSIX I/O**

**Memory allocation/deallocation routines**

**Wrapper library for an external package**

No modification to the binary executable!

Enable other TAU options (communication matrix, OTF2, event-based sampling)

# TAU Hands-On: Quickstart Guide

## Setup:

- % module load tau ; tar zxf /soft/perf-tools/tau/workshop.tgz (Theta/ACLF)

## Profiling with an un-instrumented application:

MPI: % aprun -n 64 tau\_exec -T <tag> -ebs ./a.out

- MPI+OpenMP with Intel/Clang compilers: % export TAU\_OMPT\_SUPPORT\_LEVEL=full;  
% aprun -n 64 tau\_exec -T ompt,v5 -ompt ./a.out
- Pthread: % aprun -n 64 tau\_exec -T mpi,pthread -ebs ./a.out
- Python+MPI+Sampling: % aprun -n 64 tau\_python -ebs ./a.py
- Python+MPI+CUDA+Sampling: % aprun -np 64 tau\_python -cuhti -ebs ./a.py
- C+CUDA (no MPI): % tau\_exec -T cuhti,serial -cuhti ./a.out

Analysis: % pprof -a -m | more; % paraprof (GUI)

## Tracing:

- Vampir: MPI: % export TAU\_TRACE=1; export TAU\_TRACE\_FORMAT=otf2  
% aprun -n 64 tau\_exec ./a.out; vampir traces.otf2 &
- Chrome: % export TAU\_TRACE=1; aprun -np 64 tau\_exec ./a.out; tau\_treemerge.pl;  
% tau\_trace2json tau.trc tau.edf -chrome -ignoreatomic -o app.json  
Chrome browser: chrome://tracing (Load -> app.json)
- Jumpshot: % export TAU\_TRACE=1; aprun -np 64 tau\_exec ./a.out; tau\_treemerge.pl;  
% tau2slog2 tau.trc tau.edf -o app.slog2; jumpshot app.slog2 &

# Simplifying TAU's usage (tau\_exec)

Uninstrumented execution

```
% aprun -n 64 ./a.out
```

Track MPI performance

```
% aprun -n 64 tau_exec ./a.out
```

Use event based sampling (compile with -g)

```
% aprun -n 64 tau_exec -ebs ./a.out
```

Also -ebs\_source=<PAPI\_COUNTER> -ebs\_period=<overflow\_count> -ebs\_resolution=[file|function|line]  
(line is default)

Track POSIX I/O and MPI performance (MPI enabled by default)

```
% aprun -n 64 tau_exec -T mpi,pdt,papi -io ./a.out
```

Track OpenMP runtime routines

```
% aprun -n 64 tau_exec -T ompt,v5,pdt,mpi -ompt ./a.out
```

Track memory operations

```
% export TAU_TRACK_MEMORY_LEAKS=1
```

```
% aprun -n 64 tau_exec -memory_debug ./a.out (bounds check)
```

Load wrapper interposition library

```
% aprun -n 64 tau_exec -loadlib=<path/libwrapper.so> ./a.out
```

# RUNTIME PRELOADING

- Injects TAU DSO in the executing application
- Requires dynamic executables
- We must compile with `-dynamic -g`
- Use `tau_exec` while launching the application

# Copy the workshop tarball

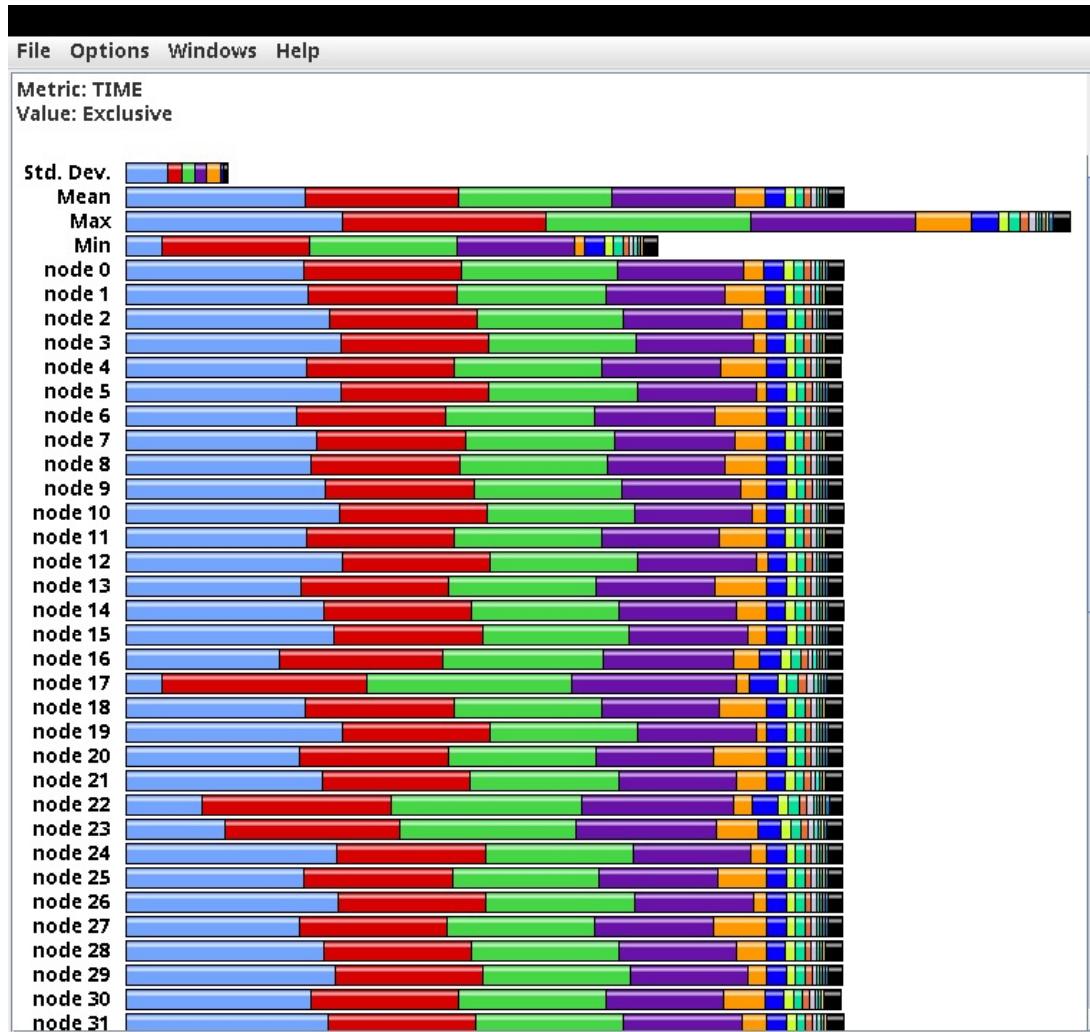
Setup preferred program environment compilers

Default set Intel Compilers with Intel MPI. You must compile with **-dynamic -g**

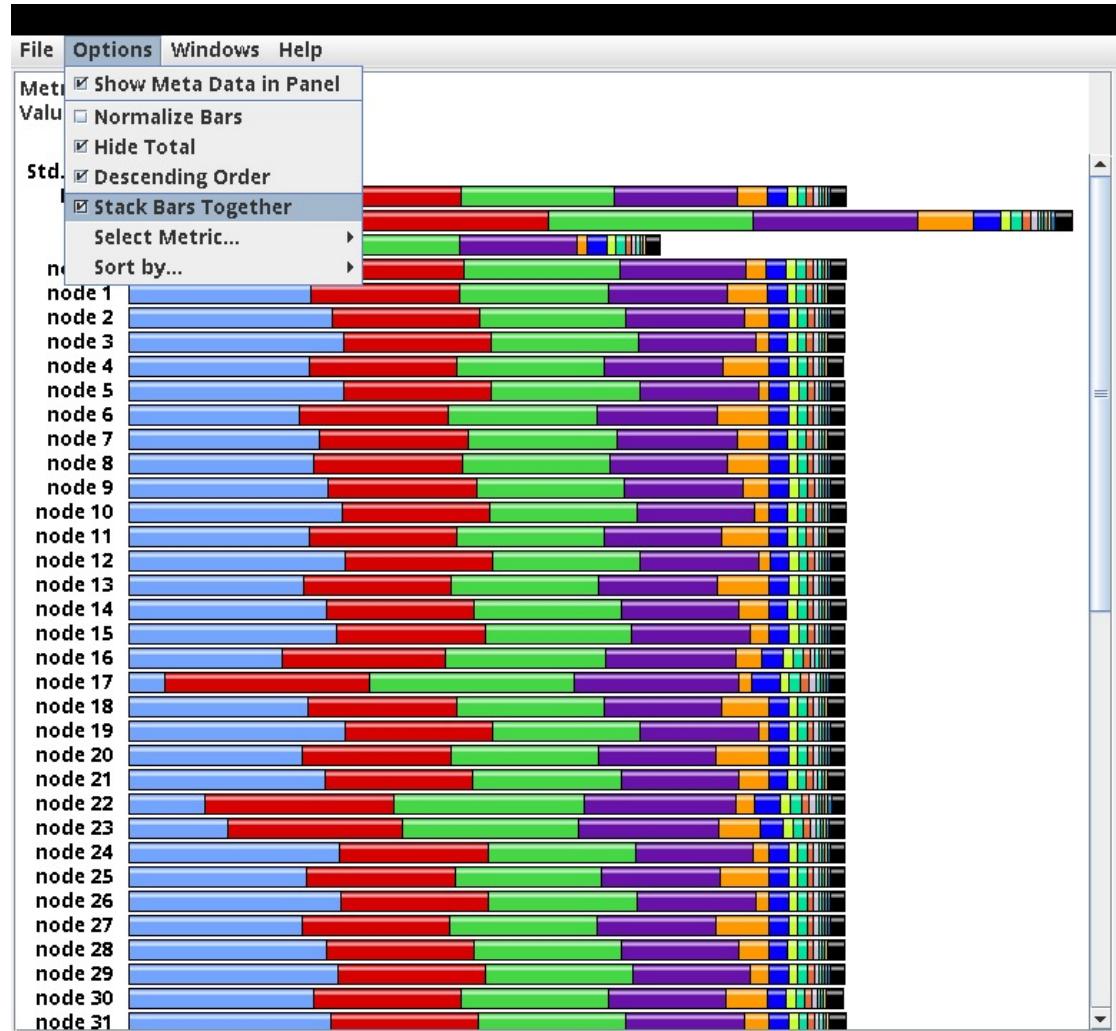
```
% module load tau;
% tar zxf /soft/perf-tools/tau/workshop.tgz
% cd workshop/MZ-NPB3.3-MPI; cat README
% make clean
% make suite
% cd bin
In a second window:
% qsub -I -n 1 -A ATPESC_Day_8 -q ATPESC2020 -t 50
% cd bin; module load tau
% export OMP_NUM_THREADS=4
% aprun -n 16 ./bt-mz.B.16 (or ./r)
% export TAU_OMPT_SUPPORT_LEVEL=full
% aprun -n 16 tau_exec -T ompt,v5,mpi,pdt -ompt ./bt-mz.B.16
(OR ./rt)
% paraprof --pack ex1.ppk
In the first window or scp to laptop and launch locally:
% paraprof ex1.ppk &
```

# ParaProf Profile Browser

% paraprof

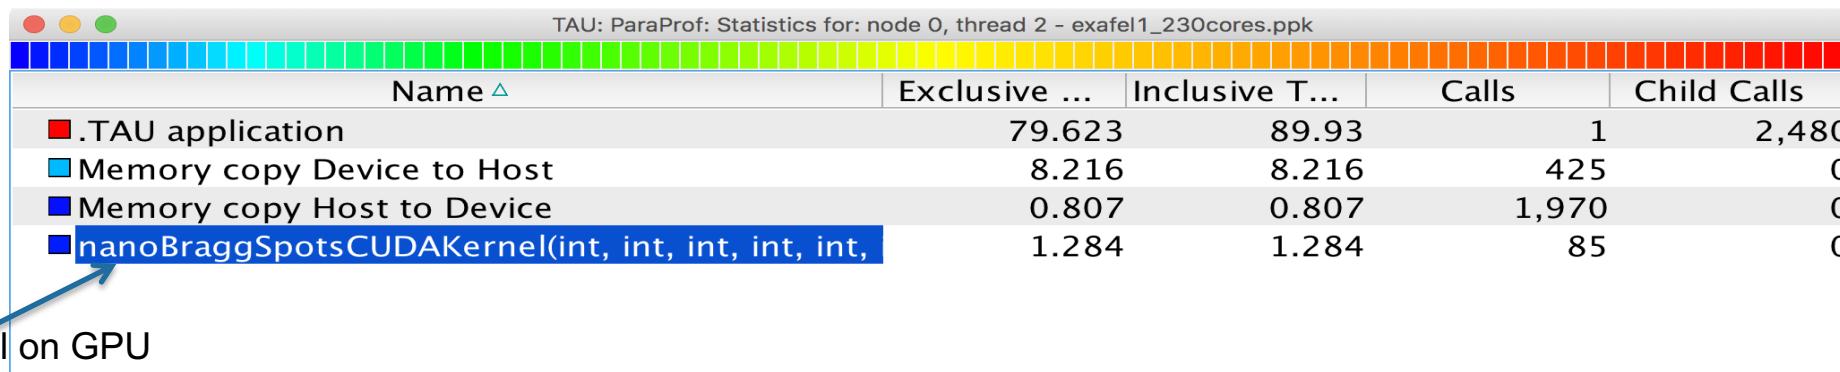


# ParaProf Profile Browser



# TAU supports Python, MPI, and CUDA

Without any modification to the source code or DSOs or interpreter, it instruments and samples the application using Python, MPI, and CUDA instrumentation.



```
% aprun -np 230 tau_python -T cupti,mpi,pdt -ebs -cupti ./exafel.py
```

Instead of:

```
% aprun -np 230 python ./exafel.py
```

# TAU Thread Statistics Table

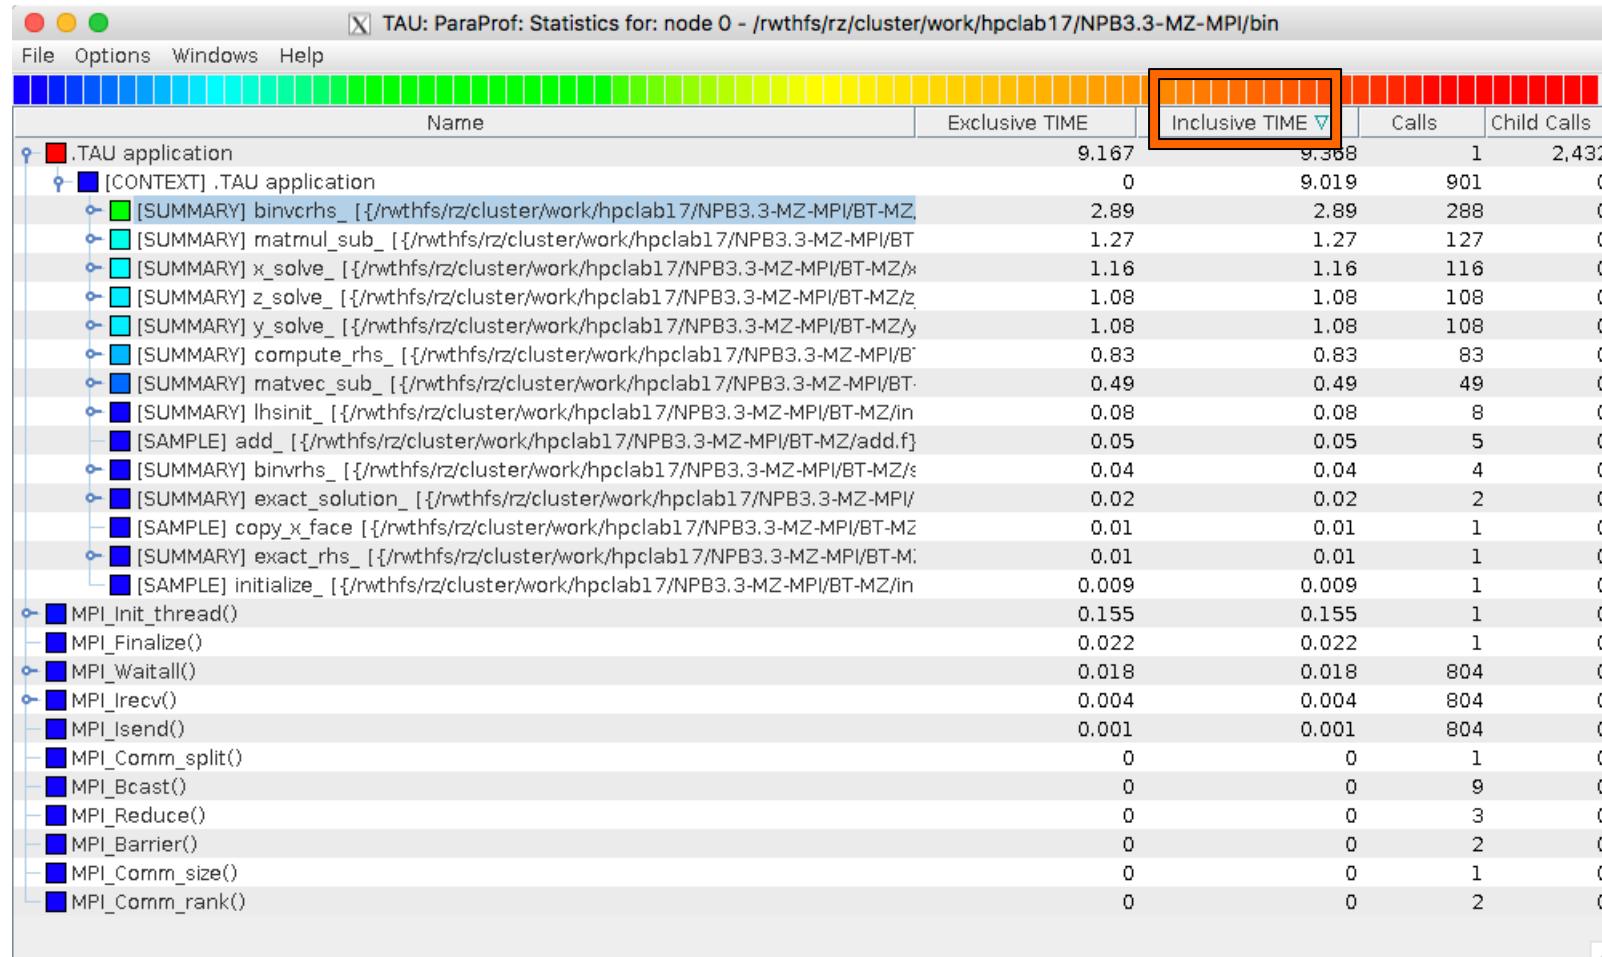
Name	Exclusive...	Inclusive ...	Calls	Child Calls
► __init__ [{from_scatterers_fft.py}{13}]	20.036	20.362	303	10,914
► run_sim2smv [{step5_pad.py}{138}]	16.78	134.9	1	1,066
► __init__ [{__init__.py}{150}]	11.669	15.909	101	1,010
► channel_pixels [{step5_pad.py}{79}]	11.029	107.657	100	13,358
► [CONTEXT] channel_pixels [{step5_pad.py}{79}]	0	9.345	312	0
█ [SAMPLE] nanoBraggSpotsCUDA [/autofs/nccs-svm1_home1/iris/adse13_161/psana-legion/simtbx/sun]	4.755	4.755	159	0
█ [SAMPLE] simtbx::nanoBragg::nanoBragg::add_nanoBragg_spots_cuda() [/autofs/nccs-svm1_home1/iris/]	4.08	4.08	136	0
█ [SAMPLE] __memset_power8 [{ }{0}]	0.3	0.3	10	0
█ [SAMPLE] UNRESOLVED /usr/lib64/libc-2.17.so	0.181	0.181	6	0
► █ [SUMMARY] Tau_handle_driver_api_memcpy(void*, CUpti_CallbackDomain, unsigned int, CUpti_CallbackDa	0.03	0.03	1	0
► cuMemcpyDtoH_v2	9.483	9.483	500	0
► expand_to_p1_iselection [{__init__.py}{1376}]	7.349	7.35	101	606
► load	7.004	7.009	2	2,251
► fp_fdp_at_wavelength [{fdp_plot.py}{44}]	6.197	6.553	100	47,550
► is_unique_set_under_symmetry [{__init__.py}{790}]	5.913	5.915	202	808
► __import__	5.782	15.766	382	78
► MPI_Init_thread()	5.616	5.723	800	1,600
► cuDevicePrimaryCtxRetain	4.987	4.987	1	0
► <module> [{__init__.py}{1}]	4.735	4.735	2	0
► MPI_Finalize()	4.255	23.888	85	756
► match_bijvoet_mates [{__init__.py}{1032}]	3.829	3.829	1	1
► bcast	3.146	3.684	101	707
► __init__ [{__init__.py}{20}]	3.073	3.448	1	9
► compute_f_mask [{__init__.py}{299}]	3.011	3.399	101	149,196
	2.897	18.853	101	707

Python, MPI, CUDA, and samples from DSOs are all integrated in a single view

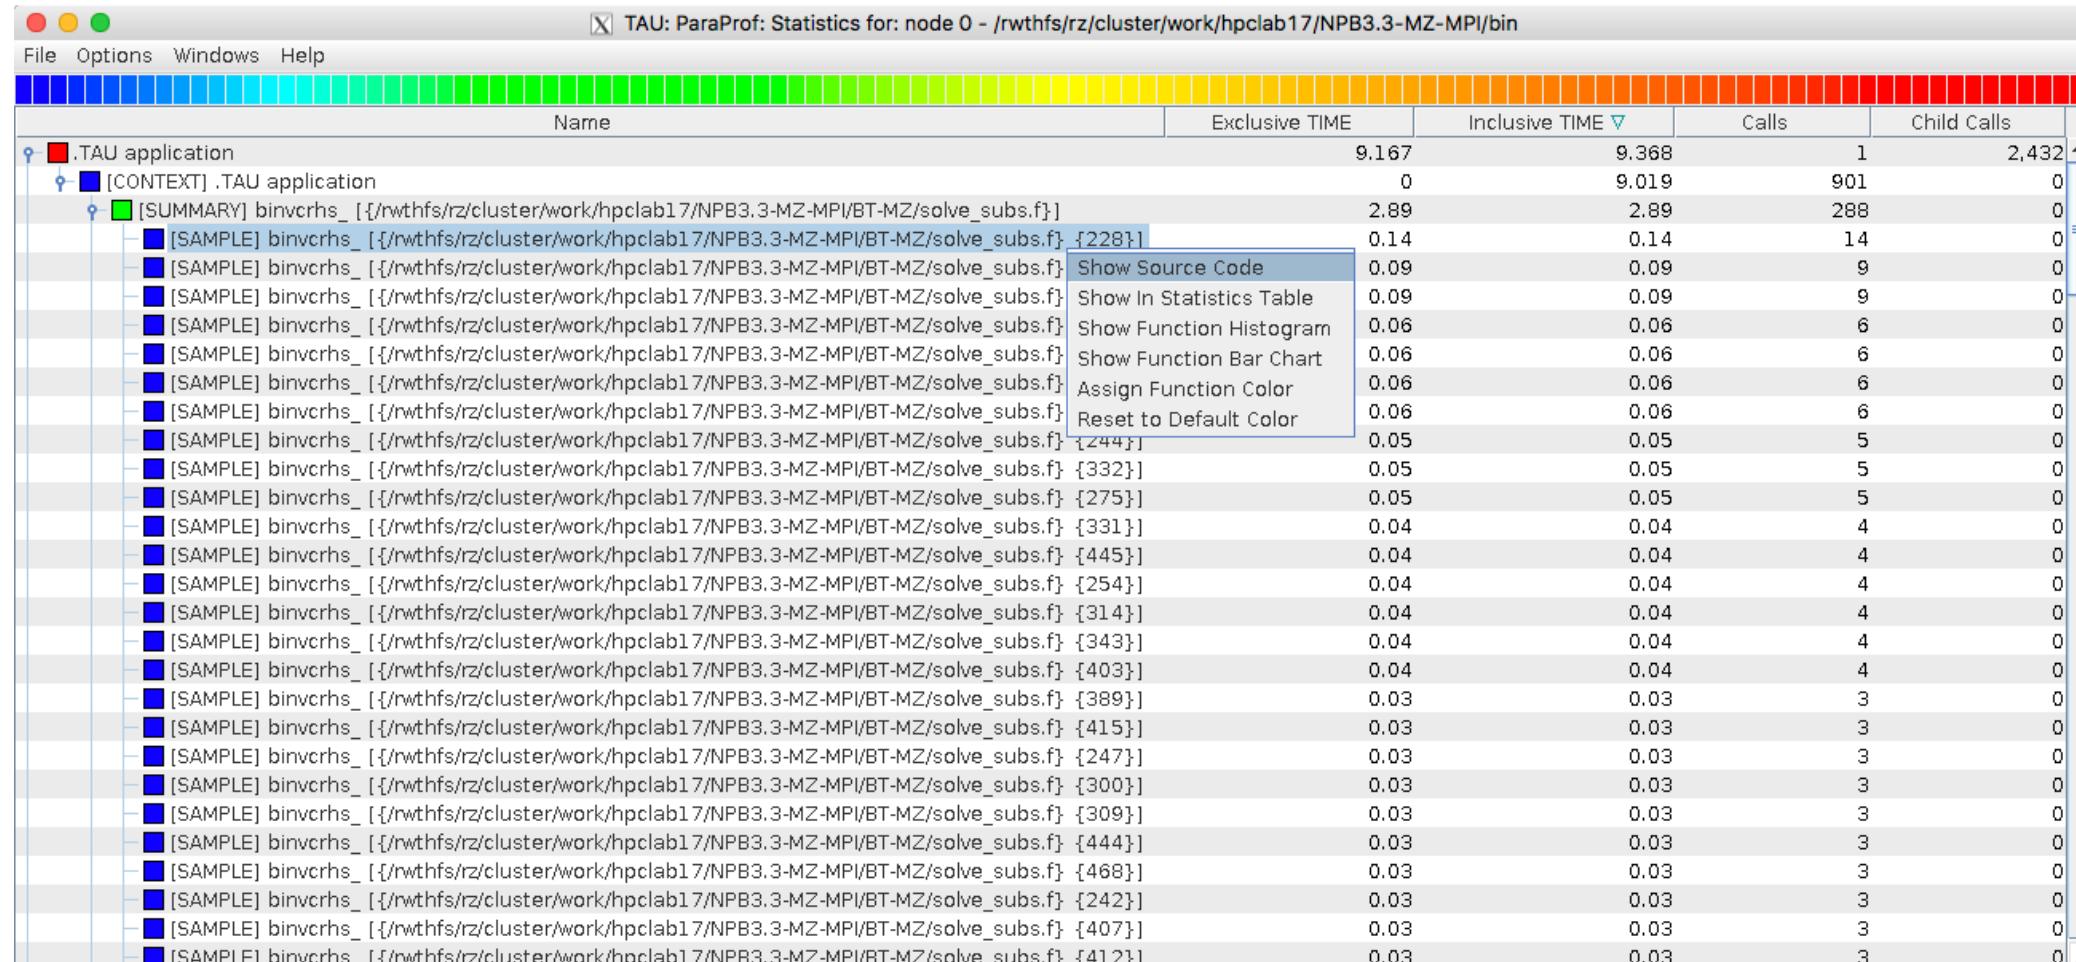
# ParaProf

Click on Columns:  
to sort by incl time

Open binvcrhs  
Click on Sample



# ParaProf



# TAU Context Event Window

Name ▲	Total	NumSamples	MaxValue	MinValue	MeanValue	Std. Dev.
<module> [{step5_batch.py}{1}]						
tst_one [{step5_batch.py}{23}]						
run_sim2smv [{step5_pad.py}{138}]						
channel_pixels [{step5_pad.py}{79}]						
cudaMemcpy						
Bytes copied from Device to Host	15,300,000,000	500	36,000,000	9,000,000	30,600,000	10,800,000
Bytes copied from Host to Device	15,423,816,000	2,300	36,000,000	8	6,706,006.957	13,564,989.185
cuMemcpyHtoD_v2						
Bytes copied from Host to Device	15,423,816,000	2,300	36,000,000	8	6,706,006.957	13,564,989.185
cuMemcpyDtoH_v2						
Bytes copied from Device to Host	15,300,000,000	500	36,000,000	9,000,000	30,600,000	10,800,000
Bytes copied from Device to Host	30,600,000,000	1,000	36,000,000	9,000,000	30,600,000	10,800,000
Bytes copied from Host to Device	30,847,632,000	4,600	36,000,000	8	6,706,006.957	13,564,989.185
Message size for broadcast	827,971,798	2	827,971,794	4	413,985,899	413,985,895

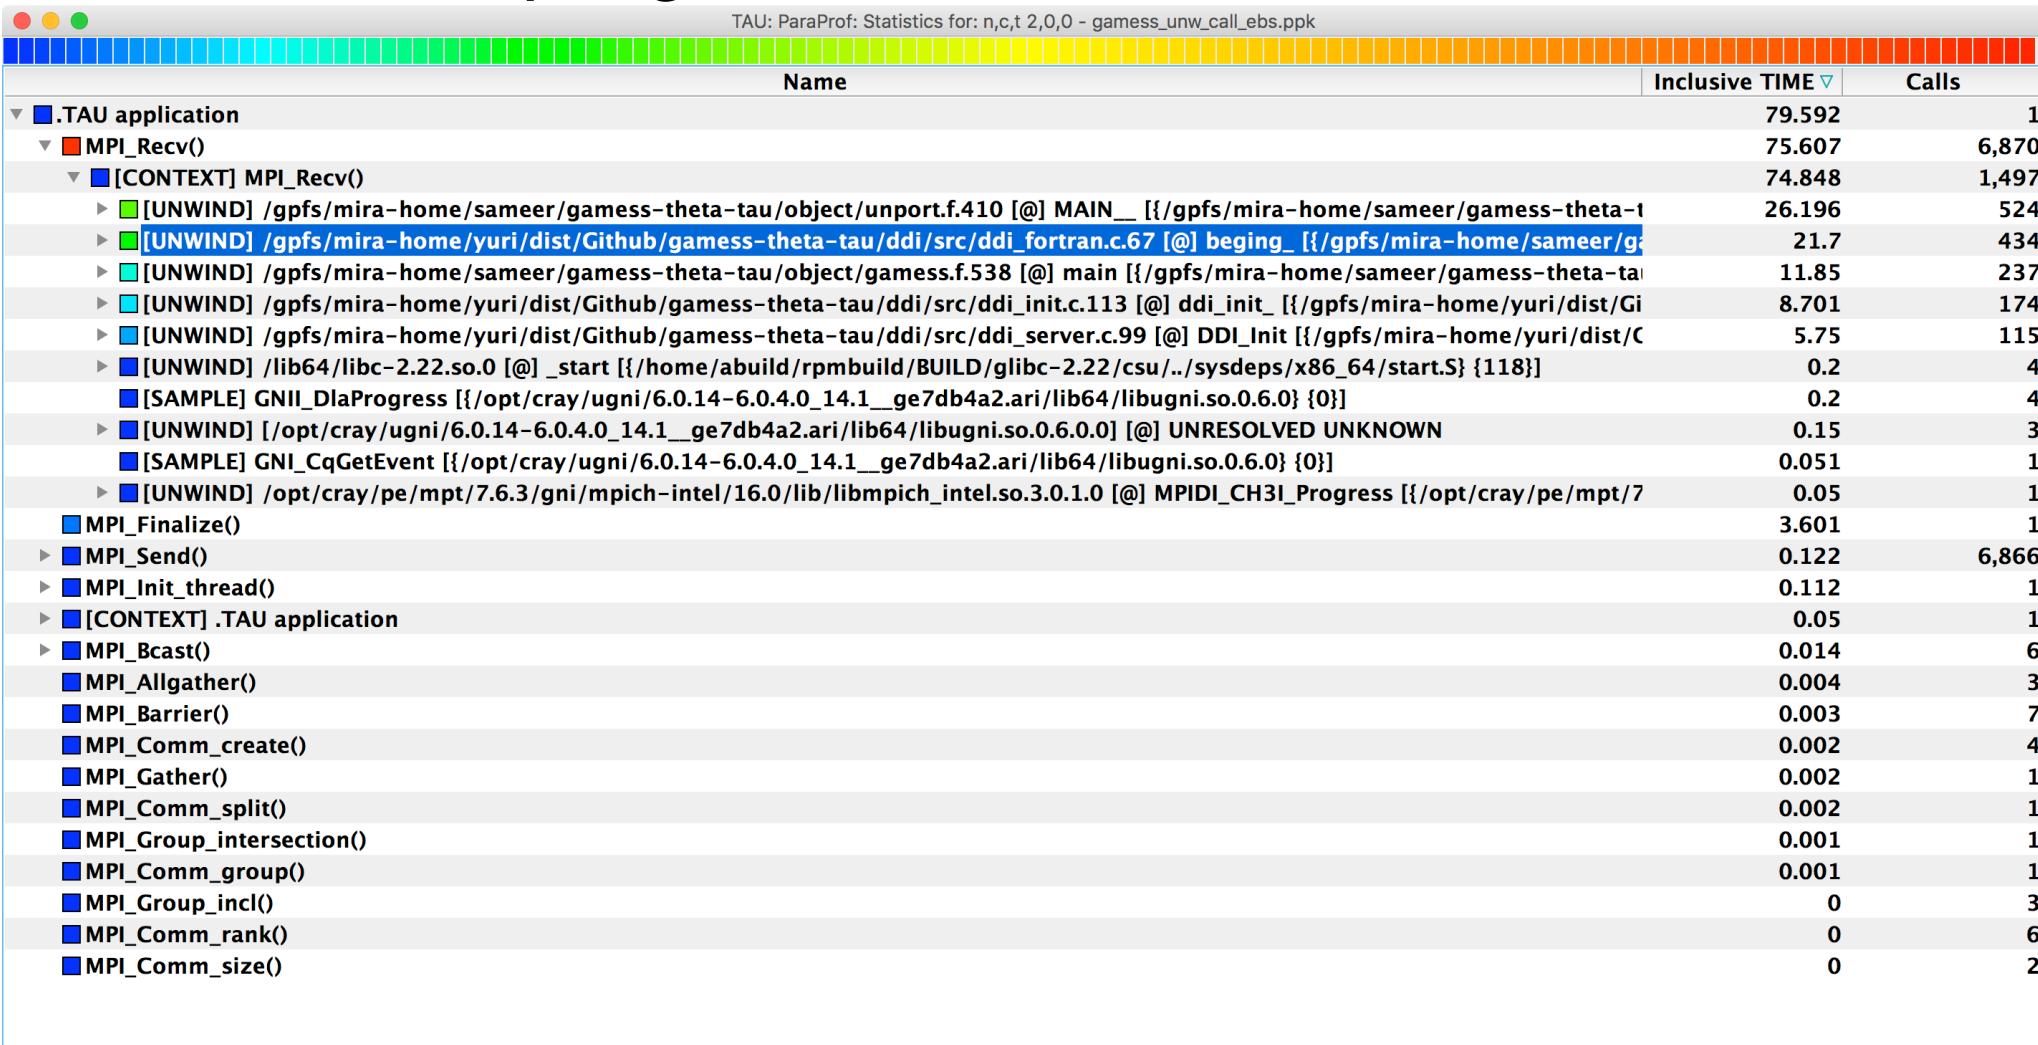
TAU tracks the data transfers between the host and the GPU.

# TAU's tracking of Python and MPI

Name	Exclusive...	Inclusive ...	Calls	Child ...
► <span style="color: cyan;">█</span> __init__ [{from_scatterers_fft.py}{13}]	19.845	20.166	303	10,914
► <span style="color: orange;">█</span> run_sim2smv [{step5_pad.py}{138}]	16.672	133.715	1	1,066
► <span style="color: blue;">█</span> MPI_Bcast()	12.263	12.263	2	0
▼ <span style="color: blue;">█</span> [CONTEXT] MPI_Bcast()	0	12.21	407	0
<span style="color: blue;">█</span> [SAMPLE] PAMI_Context_lock [{/autofs/nccs-svm1_sw/summit/.swci/1-compute/opt/spack}]	3.27	3.27	109	0
<span style="color: blue;">█</span> [SAMPLE] pthread_spin_lock [{/usr/lib64/libpthread-2.17.so}{0}]	2.34	2.34	78	0
<span style="color: blue;">█</span> [SAMPLE] start_libcoll_blocking_collective [{/autofs/nccs-svm1_sw/summit/.swci/1-compute/opt}]	1.89	1.89	63	0
<span style="color: blue;">█</span> [SAMPLE] PAMI::Device::IBV::Device::advance() [{/autofs/nccs-svm1_sw/summit/.swci/1-compute/opt}]	1.56	1.56	52	0
<span style="color: blue;">█</span> [SAMPLE] PAMI_Context_advancev [{/autofs/nccs-svm1_sw/summit/.swci/1-compute/opt}]	0.69	0.69	23	0
<span style="color: blue;">█</span> [SAMPLE] UNRESOLVED /usr/lib64/libmlx5.so.1.0.0	0.51	0.51	17	0
▼ <span style="color: blue;">█</span> [SUMMARY] LIBCOLL_Advance_pami [{/_SMPI_build_dir_____ibmsrc/r}]	0.42	0.42	14	0
<span style="color: blue;">█</span> [SAMPLE] LIBCOLL_Advance_pami [{/_SMPI_build_dir_____ibmsrc/n}]	0.42	0.42	14	0
<span style="color: blue;">█</span> [SAMPLE] PAMI_Context_unlock [{/autofs/nccs-svm1_sw/summit/.swci/1-compute/opt/s}]	0.39	0.39	13	0
<span style="color: blue;">█</span> [SAMPLE] pthread_spin_unlock [{/usr/lib64/libpthread-2.17.so}{0}]	0.36	0.36	12	0
<span style="color: blue;">█</span> [SAMPLE] __memcpy_power7 [{}{0}]	0.33	0.33	11	0
<span style="color: blue;">█</span> [SAMPLE] 0000003d.plt_call.PAMI_Context_lock [{}{0}]	0.15	0.15	5	0
<span style="color: blue;">█</span> [SAMPLE] verbs_get_exp_ctx [{pami.cc}{0}]	0.09	0.09	3	0
<span style="color: blue;">█</span> [SAMPLE] PAMI_Context_trylock_advancev [{/autofs/nccs-svm1_sw/summit/.swci/1-compute/opt}]	0.06	0.06	2	0
<span style="color: blue;">█</span> [SAMPLE] 0000003d.plt_call.PAMI_Context_unlock [{}{0}]	0.06	0.06	2	0
<span style="color: blue;">█</span> [SAMPLE] opal_progress [{/autofs/nccs-svm1_sw/summit/.swci/1-compute/opt/spack/20}]	0.03	0.03	1	0
<span style="color: blue;">█</span> [SAMPLE] 00000052.plt_call.PAMI_Context_advancev [{}{0}]	0.03	0.03	1	0
▼ <span style="color: blue;">█</span> [SUMMARY] CCMI::Executor::ShmemBroadcastT<false, CCMI::Executor::ShmemAtomicBarrier>	0.03	0.03	1	0
<span style="color: blue;">█</span> [SAMPLE] CCMI::Executor::ShmemBroadcastT<false, CCMI::Executor::ShmemAtomicBarrier>	0.03	0.03	1	0
► <span style="color: blue;">█</span> __init__ [{__init__.py}{150}]	11.518	15.698	101	1,010
► <span style="color: yellow;">█</span> channel_pixels [{step5_pad.py}{79}]	10.949	106.61	100	13,358
► <span style="color: blue;">█</span> cuMemcpyDtoH_v2	9.433	9.433	500	0

TAU can observe events in closed-source vendor libraries (e.g., in MPI\_Bcast)!

# Callstack Sampling in TAU



```
% export TAU_SAMPLING=1; export TAU_EBS_UNWIND=1
```

# UNWINDING CALLSTACKS

TAU: ParaProf: Statistics for: n,c,t 2,0,0 - gamess\_unw\_call\_ebs.ppk

	Name	Inclusive TIME ▼	Calls
▼ .TAU application		79.592	1
▼ MPI_Recv()		75.607	6,870
▼ [CONTEXT] MPI_Recv()		74.848	1,497
► [UNWIND] /gpfs/mira-home/sameer/gamess-theta-tau/object/unport.f.410 [@] MAIN__ [{/gpfs/mira-home/sameer/gamess-theta-tau}]	26.196	524	
▼ [UNWIND] /gpfs/mira-home/yuri/dist/Github/gamess-theta-tau/ddi/src/ddi_fortran.c.67 [@] begin_ [{/gpfs/mira-home/sameer/gamess-theta-tau}]	21.7	434	
▼ [UNWIND] /gpfs/mira-home/yuri/dist/Github/gamess-theta-tau/ddi/src/ddi_init.c.113 [@] ddi_init_ [{/gpfs/mira-home/yuri/dist/Github/gamess-theta-tau}]	21.7	434	
▼ [UNWIND] /gpfs/mira-home/yuri/dist/Github/gamess-theta-tau/ddi/src/ddi_server.c.99 [@] DDI_Init [{/gpfs/mira-home/yuri/dist/Github/gamess-theta-tau}]	21.7	434	
▼ [UNWIND] /gpfs/mira-home/yuri/dist/Github/gamess-theta-tau/ddi/src/ddi_recv.c.65 [@] DDI_Server [{/gpfs/mira-home/yuri/dist/Github/gamess-theta-tau}]	21.7	434	
▼ [UNWIND] /lus/theta-fs0/software/perftools/tau/tau-2.26.3/src/Profile/TauMpi.c.2371 [@] DDI_Recv_request [{/gpfs/mira-home/yuri/dist/Github/gamess-theta-tau}]	21.7	434	
▼ [UNWIND] /opt/cray/pe/mpt/7.6.3/gni/mpich-intel/16.0/lib/libmpich_intel.so.3.0.1.0 [@] MPI_Recv [{/lus/theta-fs0/software/perftools/tau/tau-2.26.3/src/Profile/TauMpi.c.2371}]	21.7	434	
▼ [UNWIND] /opt/cray/pe/mpt/7.6.3/gni/mpich-intel/16.0/lib/libmpich_intel.so.3.0.1.0 [@] PMPI_Recv [{/opt/cray/pe/mpt/7.6.3/gni/mpich-intel/16.0/lib/libmpich_intel.so.3.0.1.0}]	21.7	434	
▼ [UNWIND] /opt/cray/pe/mpt/7.6.3/gni/mpich-intel/16.0/lib/libmpich_intel.so.3.0.1.0 [@] MPIIDI_CH3I_Progress [{/opt/cray/pe/mpt/7.6.3/gni/mpich-intel/16.0/lib/libmpich_intel.so.3.0.1.0}]	21.45	429	
▼ [UNWIND] /opt/cray/ugni/6.0.14-6.0.4.0_14.1_ge7db4a2.ari/lib64/libugni.so.0.6.0.0 [@] MPIID_nem_gni_poll [{/opt/cray/ugni/6.0.14-6.0.4.0_14.1_ge7db4a2.ari/lib64/libugni.so.0.6.0.0}]	15.95	319	
► [SAMPLE] GNI_SmsgGetNextWTag [{/opt/cray/ugni/6.0.14-6.0.4.0_14.1_ge7db4a2.ari/lib64/libugni.so.0.6.0}]	10.349	207	
► [SAMPLE] GNI_CqGetEvent [{/opt/cray/ugni/6.0.14-6.0.4.0_14.1_ge7db4a2.ari/lib64/libugni.so.0.6.0} {0}]	5.6	112	
► [UNWIND] gni_poll.c.0 [@] MPIID_nem_gni_poll [{/opt/cray/pe/mpt/7.6.3/gni/mpich-intel/16.0/lib/libmpich_intel.so.3.0.1.0}]	5.25	105	
► [UNWIND] /opt/cray/pe/mpt/7.6.3/gni/mpich-intel/16.0/lib/libmpich_intel.so.3.0.1.0 [@] MPIID_nem_gni_poll [{/opt/cray/pe/mpt/7.6.3/gni/mpich-intel/16.0/lib/libmpich_intel.so.3.0.1.0}]	0.25	5	
► [UNWIND] UNRESOLVED [@] MPIIDI_CH3I_Progress [{/opt/cray/pe/mpt/7.6.3/gni/mpich-intel/16.0/lib/libmpich_intel.so.3.0.1.0}]	0.25	5	
► [UNWIND] /gpfs/mira-home/sameer/gamess-theta-tau/object/gamess.f.538 [@] main [{/gpfs/mira-home/sameer/gamess-theta-tau}]	11.85	237	
► [UNWIND] /gpfs/mira-home/yuri/dist/Github/gamess-theta-tau/ddi/src/ddi_init.c.113 [@] ddi_init_ [{/gpfs/mira-home/yuri/dist/Github/gamess-theta-tau}]	8.701	174	
► [UNWIND] /gpfs/mira-home/yuri/dist/Github/gamess-theta-tau/ddi/src/ddi_server.c.99 [@] DDI_Init [{/gpfs/mira-home/yuri/dist/Github/gamess-theta-tau}]	5.75	115	
► [UNWIND] /lib64/libc-2.22.so.0 [@] _start [{/home/abuild/rpmbuild/BUILD/glibc-2.22/csu/../sysdeps/x86_64/start.S} {118}]	0.2	4	
► [SAMPLE] GNII_DlaProgress [{/opt/cray/ugni/6.0.14-6.0.4.0_14.1_ge7db4a2.ari/lib64/libugni.so.0.6.0} {0}]	0.2	4	
► [UNWIND] [/opt/cray/ugni/6.0.14-6.0.4.0_14.1_ge7db4a2.ari/lib64/libugni.so.0.6.0.0] [@] UNRESOLVED UNKNOWN	0.15	3	
► [SAMPLE] GNI_CqGetEvent [{/opt/cray/ugni/6.0.14-6.0.4.0_14.1_ge7db4a2.ari/lib64/libugni.so.0.6.0} {0}]	0.051	1	
► [UNWIND] /opt/cray/pe/mpt/7.6.3/gni/mpich-intel/16.0/lib/libmpich_intel.so.3.0.1.0 [@] MPIIDI_CH3I_Progress [{/opt/cray/pe/mpt/7.6.3/gni/mpich-intel/16.0/lib/libmpich_intel.so.3.0.1.0}]	0.05	1	
► [MPI_Finalize()]		3.601	1
► [MPI_Send()]		0.122	6,866
► [MPI_Init_thread()]		0.112	1
► [CONTEXT] .TAU application		0.05	1

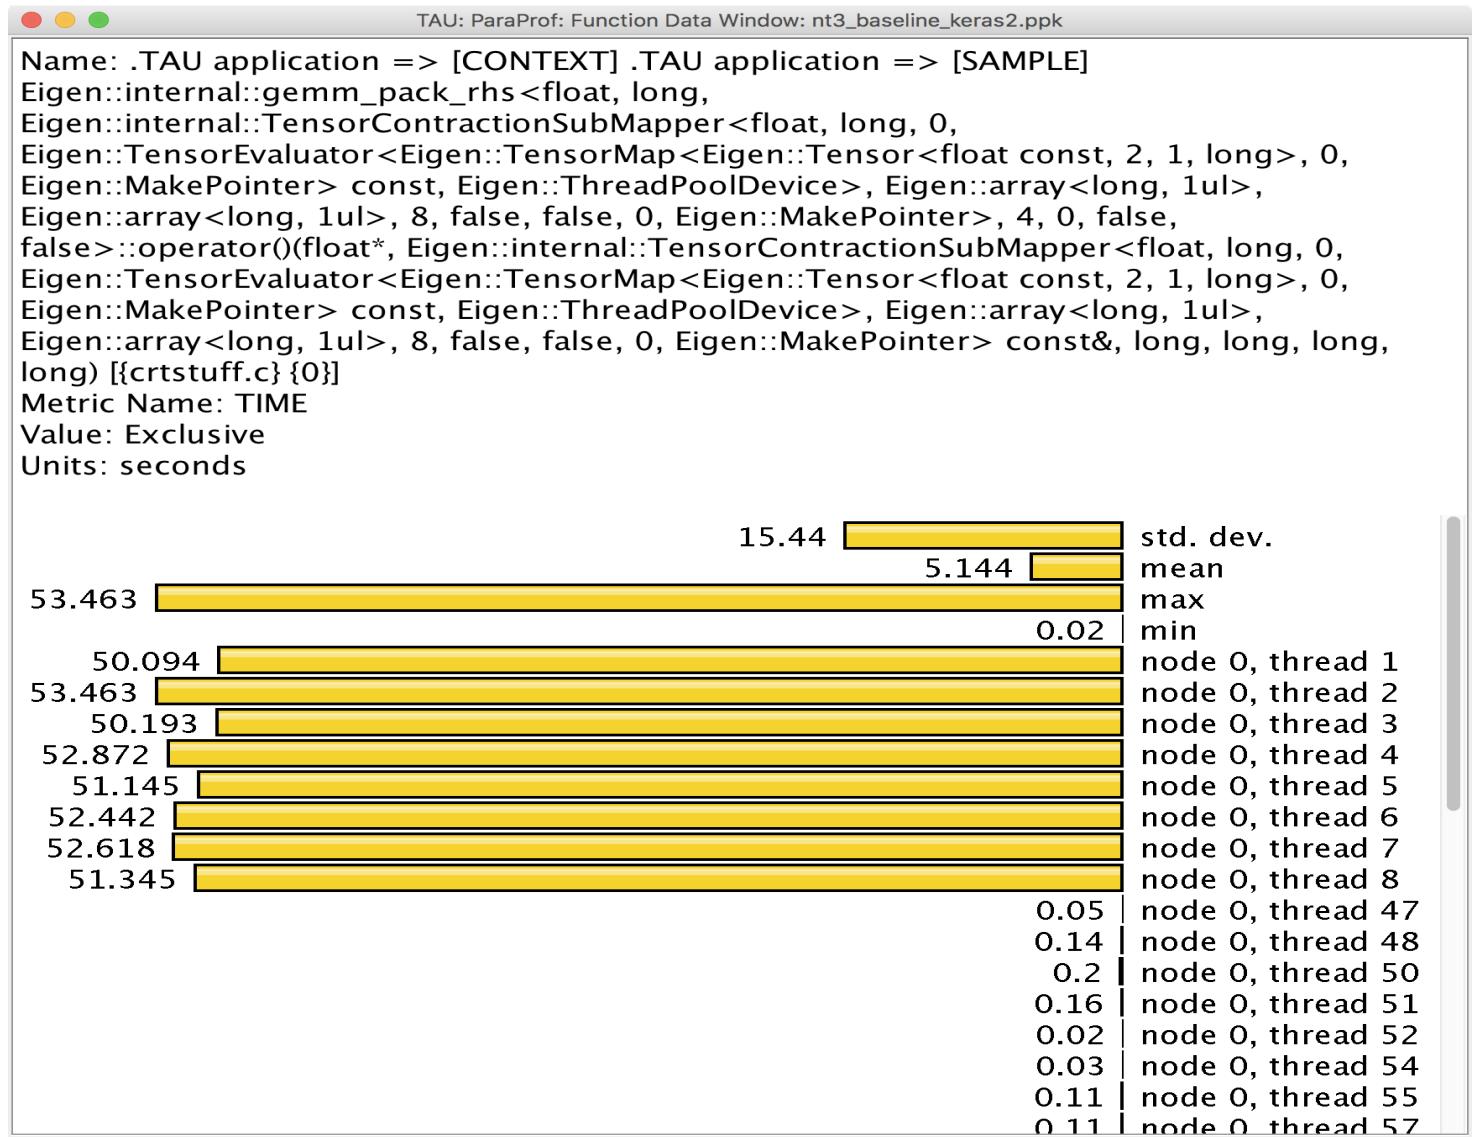
% export TAU\_SAMPLING=1; export TAU\_EBS\_UNWIND=1

# Deep Learning: Tensorflow

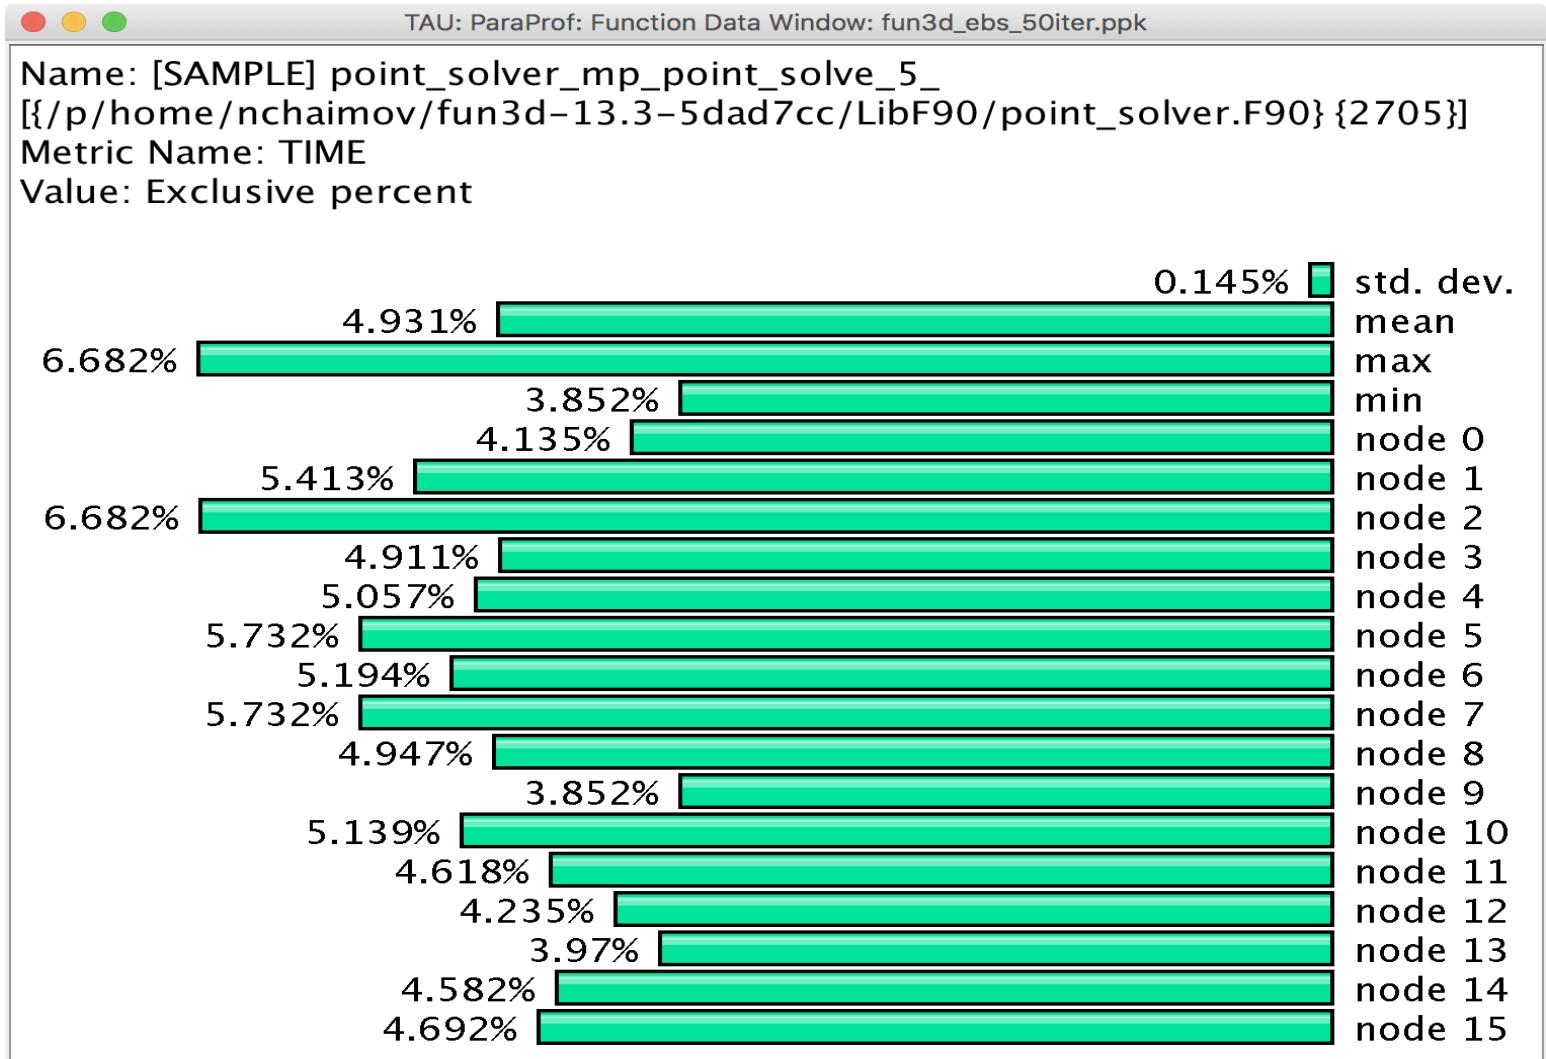
TAU: ParaProf: Statistics for: node 0, thread 8 - nt3_baseline_keras2.ppk			
	Name	Inclusiv...	Calls ▼
■ .TAU application		519.211	1
■ [CONTEXT] .TAU application		509.222	50,915
■ [SAMPLE] Eigen::internal::gebp_kernel<float, float, long, Eigen::internal::blas_data_mapper<float, long, 0, 0>,		240.632	24,089
■ [SAMPLE] __pthread_cond_wait [{ } {0}]		86.384	8,634
■ [SAMPLE] Eigen::internal::gemm_pack_rhs<float, long, Eigen::internal::TensorContractionSubMapper<float, long, 0, 0>,		51.345	5,135
■ [SAMPLE] Eigen::internal::gemm_pack_rhs<float, long, Eigen::internal::TensorContractionSubMapper<float, long, 0, 0>,		24.375	2,416
■ [SAMPLE] void tensorflow::SpatialMaxPoolWithArgMaxHelper<Eigen::ThreadPoolDevice, float>(tensorflow::OpK		16.301	1,630
■ [SAMPLE] __memset_sse2 [{ } {0}]		13.446	1,336
■ [SAMPLE] Eigen::TensorEvaluator<Eigen::TensorContractionOp<Eigen::array<Eigen::IndexPair<long>, 1ul> co		5.99	599
■ [SAMPLE] long Eigen::internal::operator/<long, false>(long const&, Eigen::internal::TensorIntDivisor<long, fals		5.843	585
■ [SAMPLE] std::__Function_handler<void (long, long), Eigen::internal::TensorExecutor<Eigen::TensorAssignOp<I		5.377	538
■ [SAMPLE] float __vector Eigen::TensorEvaluator<Eigen::TensorBroadcastingOp<Eigen::IndexList<int, Eigen::typ		4.862	487
■ [SAMPLE] Eigen::TensorEvaluator<Eigen::TensorContractionOp<Eigen::array<Eigen::IndexPair<long>, 1ul> co		4.775	478
■ [SAMPLE] Eigen::TensorEvaluator<Eigen::TensorAssignOp<Eigen::TensorMap<Eigen::Tensor<float, 1, 1, long>		4.037	404
■ [SAMPLE] Eigen::internal::gemm_pack_lhs<float, long, Eigen::internal::TensorContractionSubMapper<float, long,		3.679	367
■ [SAMPLE] Eigen::internal::EvalRange<Eigen::TensorEvaluator<Eigen::TensorAssignOp<Eigen::TensorMap<Eigen::Tens		2.981	298
■ [SAMPLE] tensorflow::MaxPoolingOp<Eigen::ThreadPoolDevice, float>::SpatialMaxPool(tensorflow::OpKernelCo		2.915	295
■ [SAMPLE] std::__Function_handler<void (long, long), Eigen::internal::TensorExecutor<Eigen::TensorAssignOp<I		2.91	291
■ [SAMPLE] std::__Function_handler<void (long, long), Eigen::internal::TensorExecutor<Eigen::TensorAssignOp<I		2.772	277
■ [SAMPLE] Eigen::internal::gemm_pack_lhs<float, long, Eigen::internal::TensorContractionSubMapper<float, long,		2.481	248
■ [SAMPLE] std::__Function_handler<void (long, long), Eigen::internal::TensorExecutor<Eigen::TensorAssignOp<I		2.148	215
■ [SAMPLE] void Eigen::internal::call_dense_assignment_loop<Eigen::Map<Eigen::Matrix<float, -1, -1, 0, -1, -1>		2.008	197
■ [SAMPLE] Eigen::NonBlockingThreadPoolTempl<tensorflow::thread::EigenEnvironment>::WorkerLoop(int) [{/ho		1.999	200
■ [SAMPLE] Eigen::internal::ptr transpose(Eigen::internal::PacketBlock<float __vector, 4>&) [{crtstuff.c} {0}]		1.919	192
■ [SAMPLE] Eigen::internal::gemm_pack_rhs<float, long, Eigen::internal::TensorContractionSubMapper<float, long,		1.607	160
■ [SAMPLE] Eigen::TensorEvaluator<Eigen::TensorContractionOp<Eigen::array<Eigen::IndexPair<long>, 1ul> co		1.518	152

% tau\_python -ebs nt3\_baseline\_keras2.py (CANDLE)

# Sampling Tensorflow

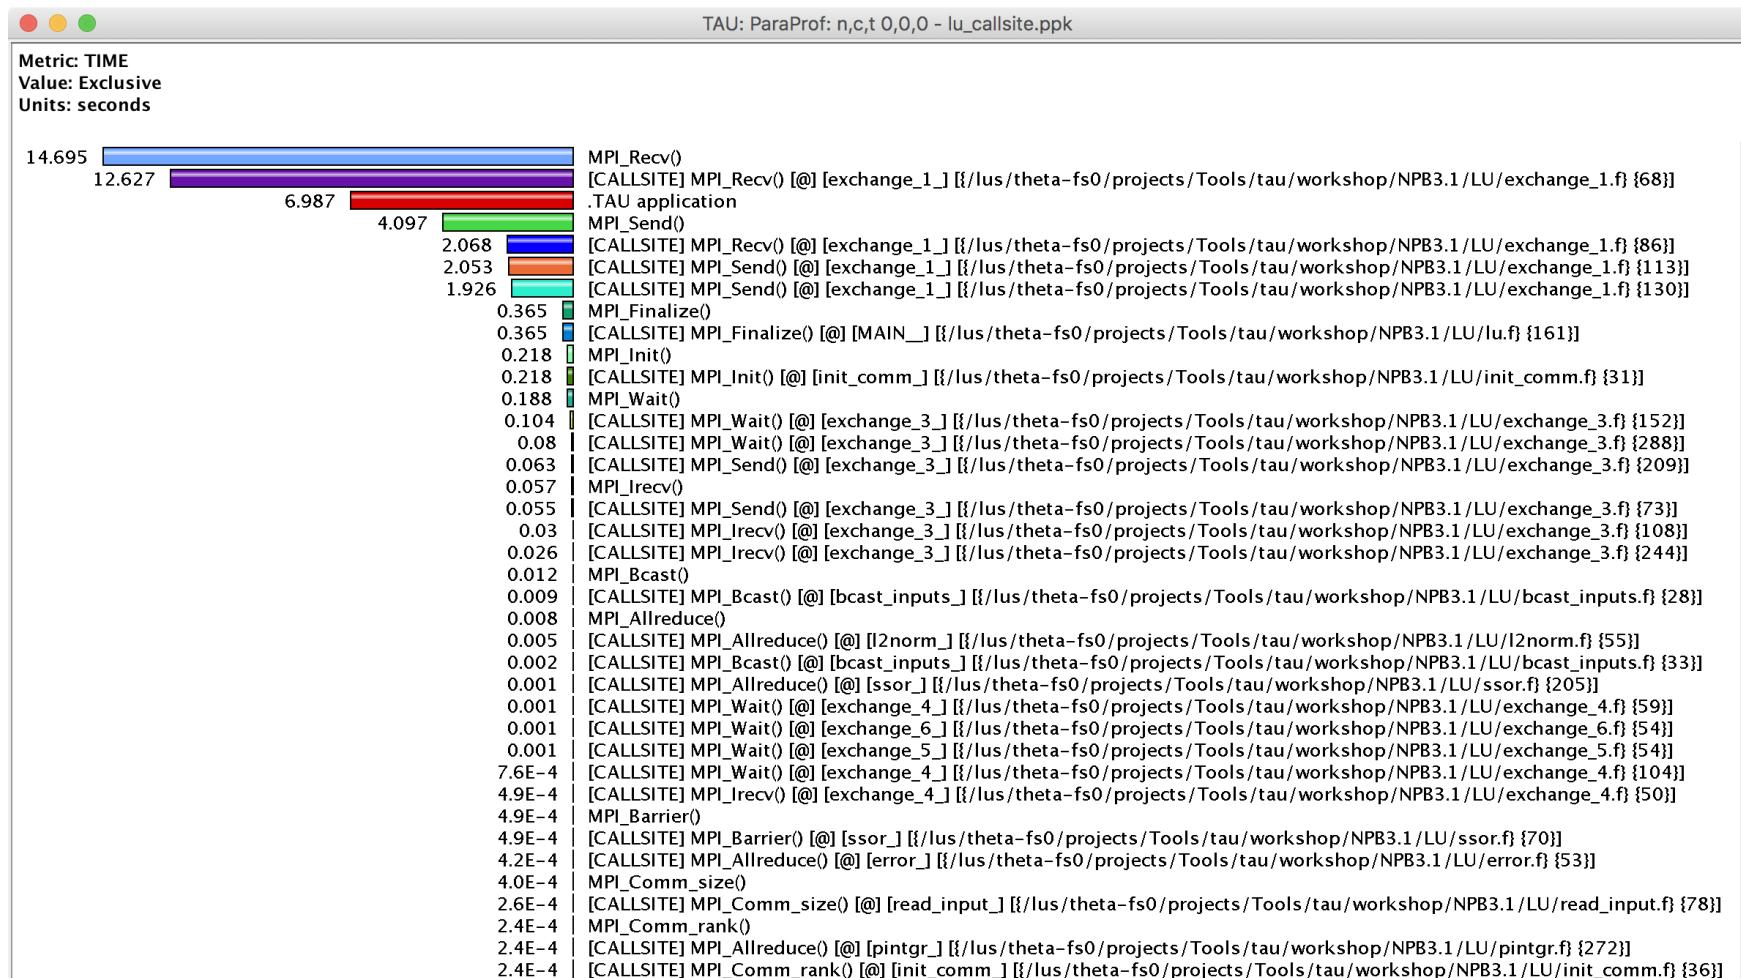


# Event Based Sampling (EBS)



```
% aprun -n 16 tau_exec -ebs a.out
```

# Callsite Profiling and Tracing



% export TAU\_CALLSITE=1

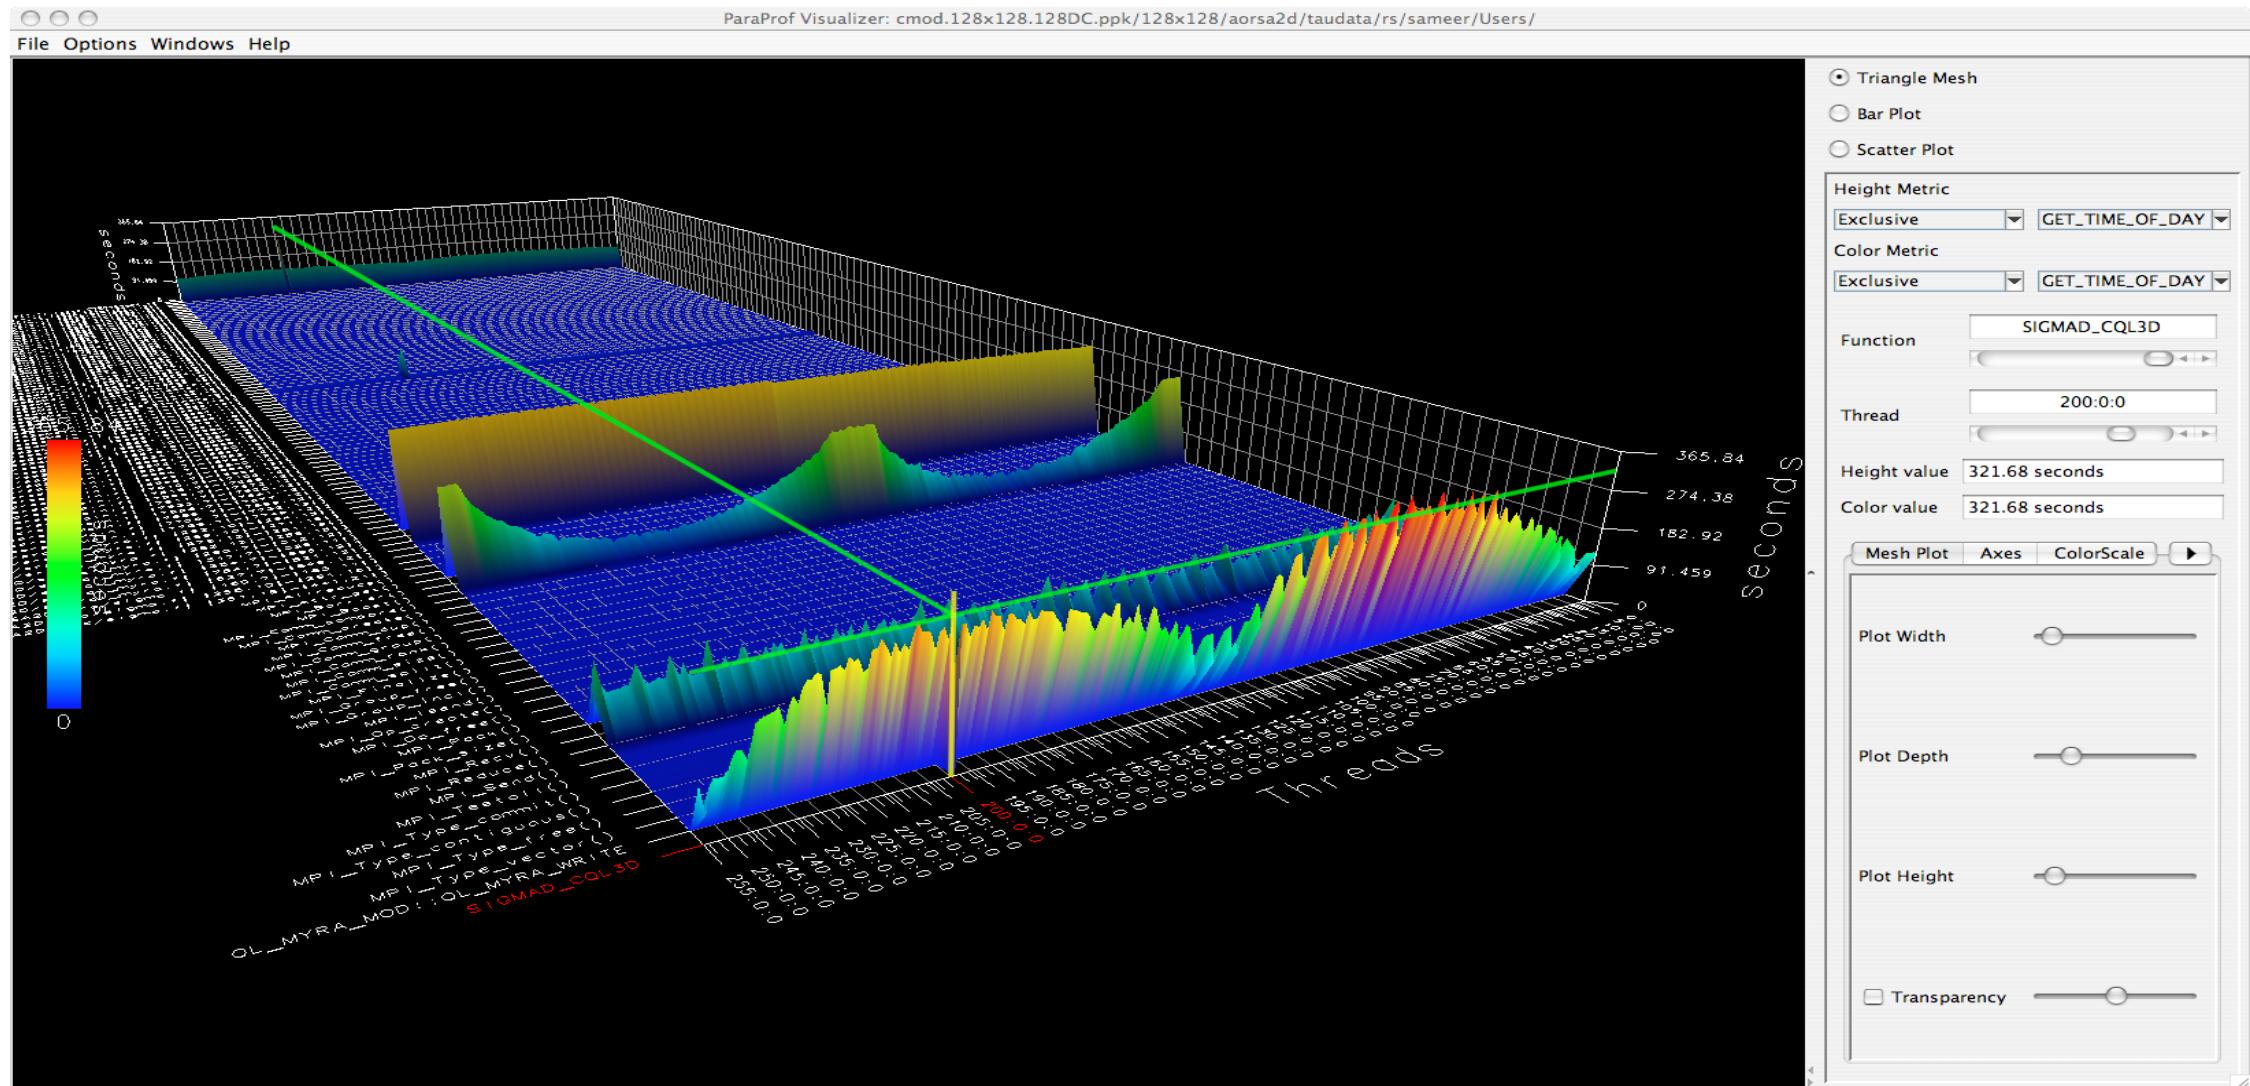
# TAU – Context Events

TAU: ParaProf: Context Events for thread: n,c,t, 1,0,0 – samarc_obe_4p_iomem_cp.ppk							
Name	Total	MeanValue	NumSamples	MinValue	MaxValue	Std. Dev.	
► .TAU application							
► read()							
► fopen64()							
► fclose()							
▼ OurMain()							
malloc size	25,235	1,097.174	23	11	12,032	2,851.143	
free size	22,707	1,746.692	13	11	12,032	3,660.642	
▼ OurMain [{wrapper.py}{3}]							
► read()							
malloc size	3,877	323.083	12	32	981	252.72	
free size	1,536	219.429	7	32	464	148.122	
► fopen64()							
► fclose()							
▼ <module> [{obe.py}{8}]							
► writeRestartData [{samarcInterface.py}{145}]							
▼ samarcWriteRestartData							
► write()							
WRITE Bandwidth (MB/s) <file="samarc/restore.00002/nodes.00004/proc.00001">	74.565	117	0	2,156.889	246.386		
WRITE Bandwidth (MB/s) <file="samarc/restore.00001/nodes.00004/proc.00001">	77.594	117	0	1,941.2	228.366		
WRITE Bandwidth (MB/s)	76.08	234	0	2,156.889	237.551		
Bytes Written <file="samarc/restore.00002/nodes.00004/proc.00001">	2,097,552	17,927.795	117	1	1,048,576	133,362.946	
Bytes Written <file="samarc/restore.00001/nodes.00004/proc.00001">	2,097,552	17,927.795	117	1	1,048,576	133,362.946	
Bytes Written	4,195,104	17,927.795	234	1	1,048,576	133,362.946	
► open64()							

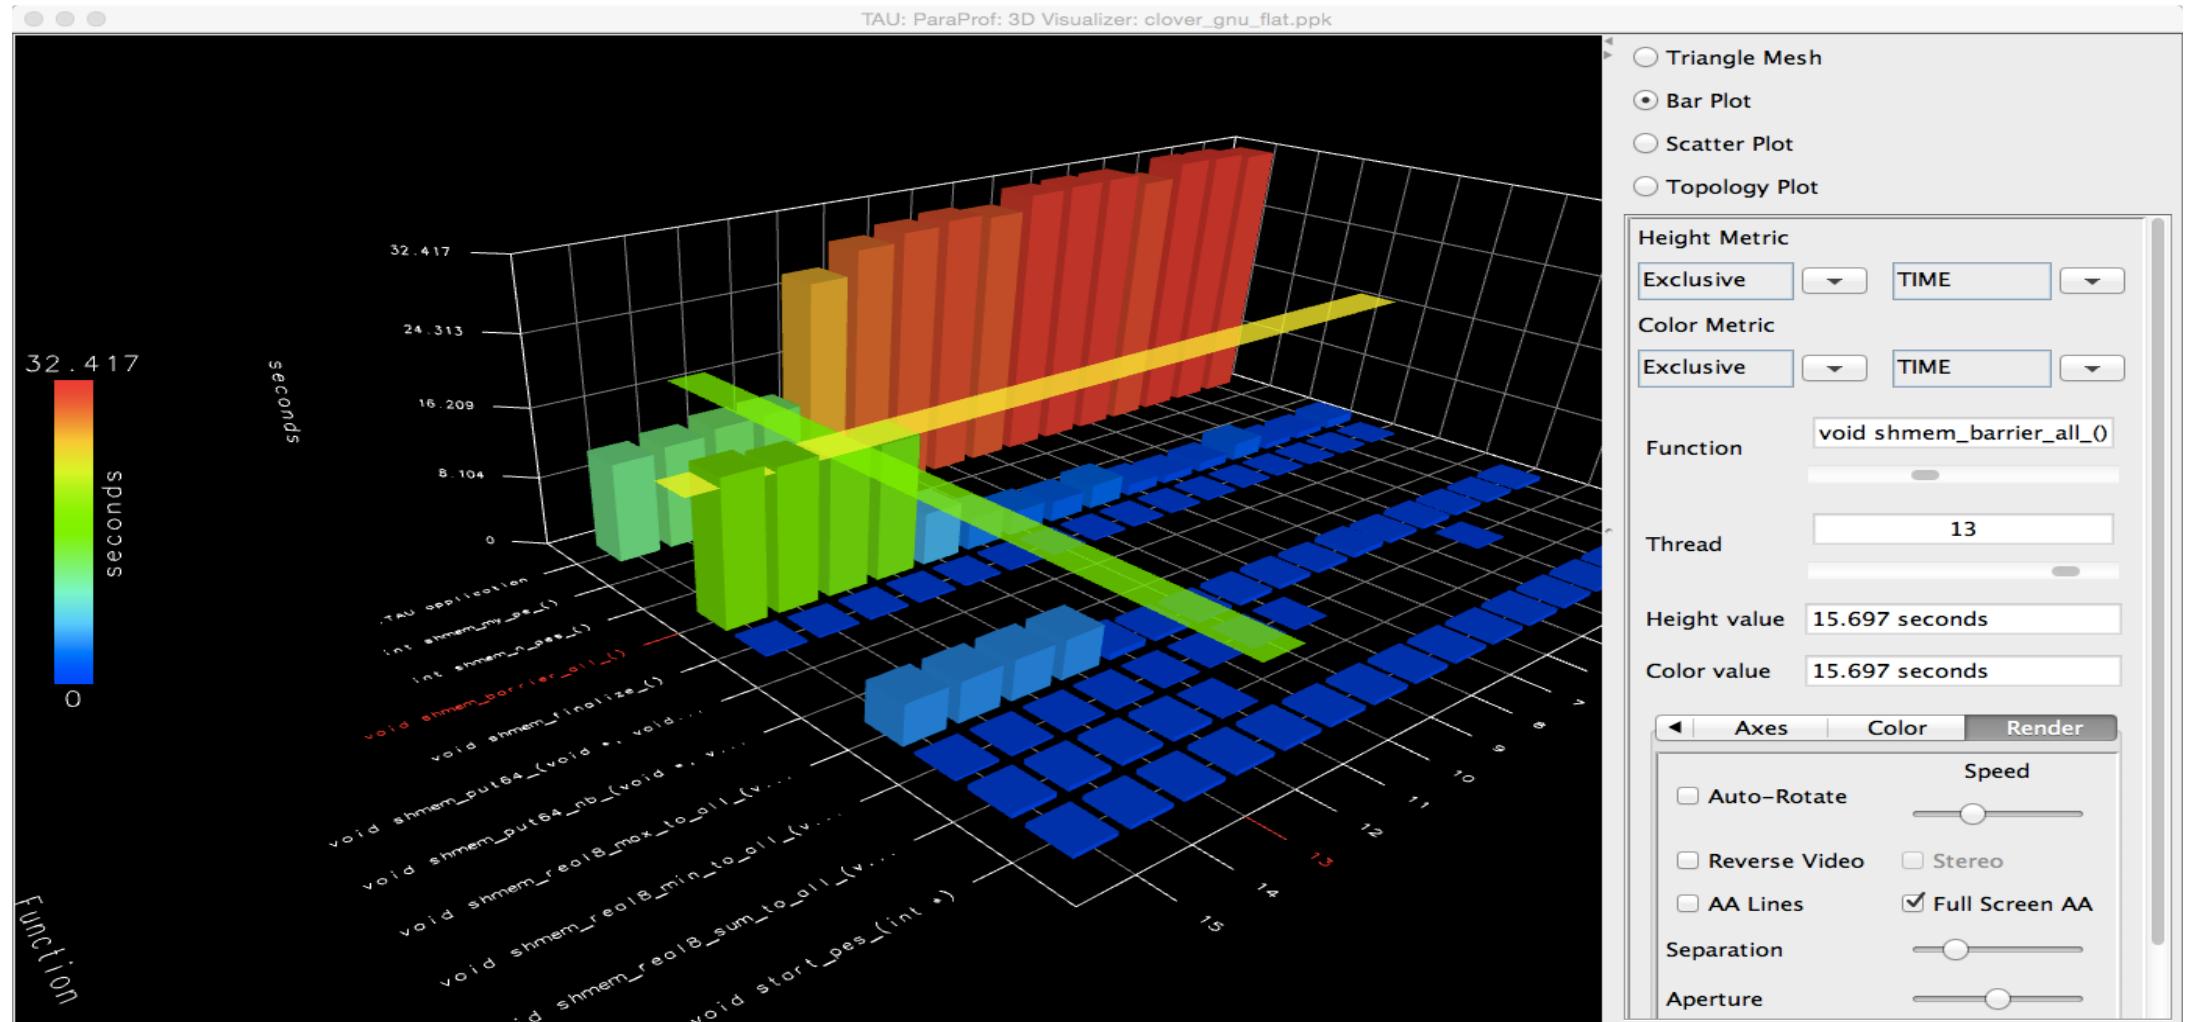
Write bandwidth per file

Bytes written to each file

# ParaProf 3D Profile Browser

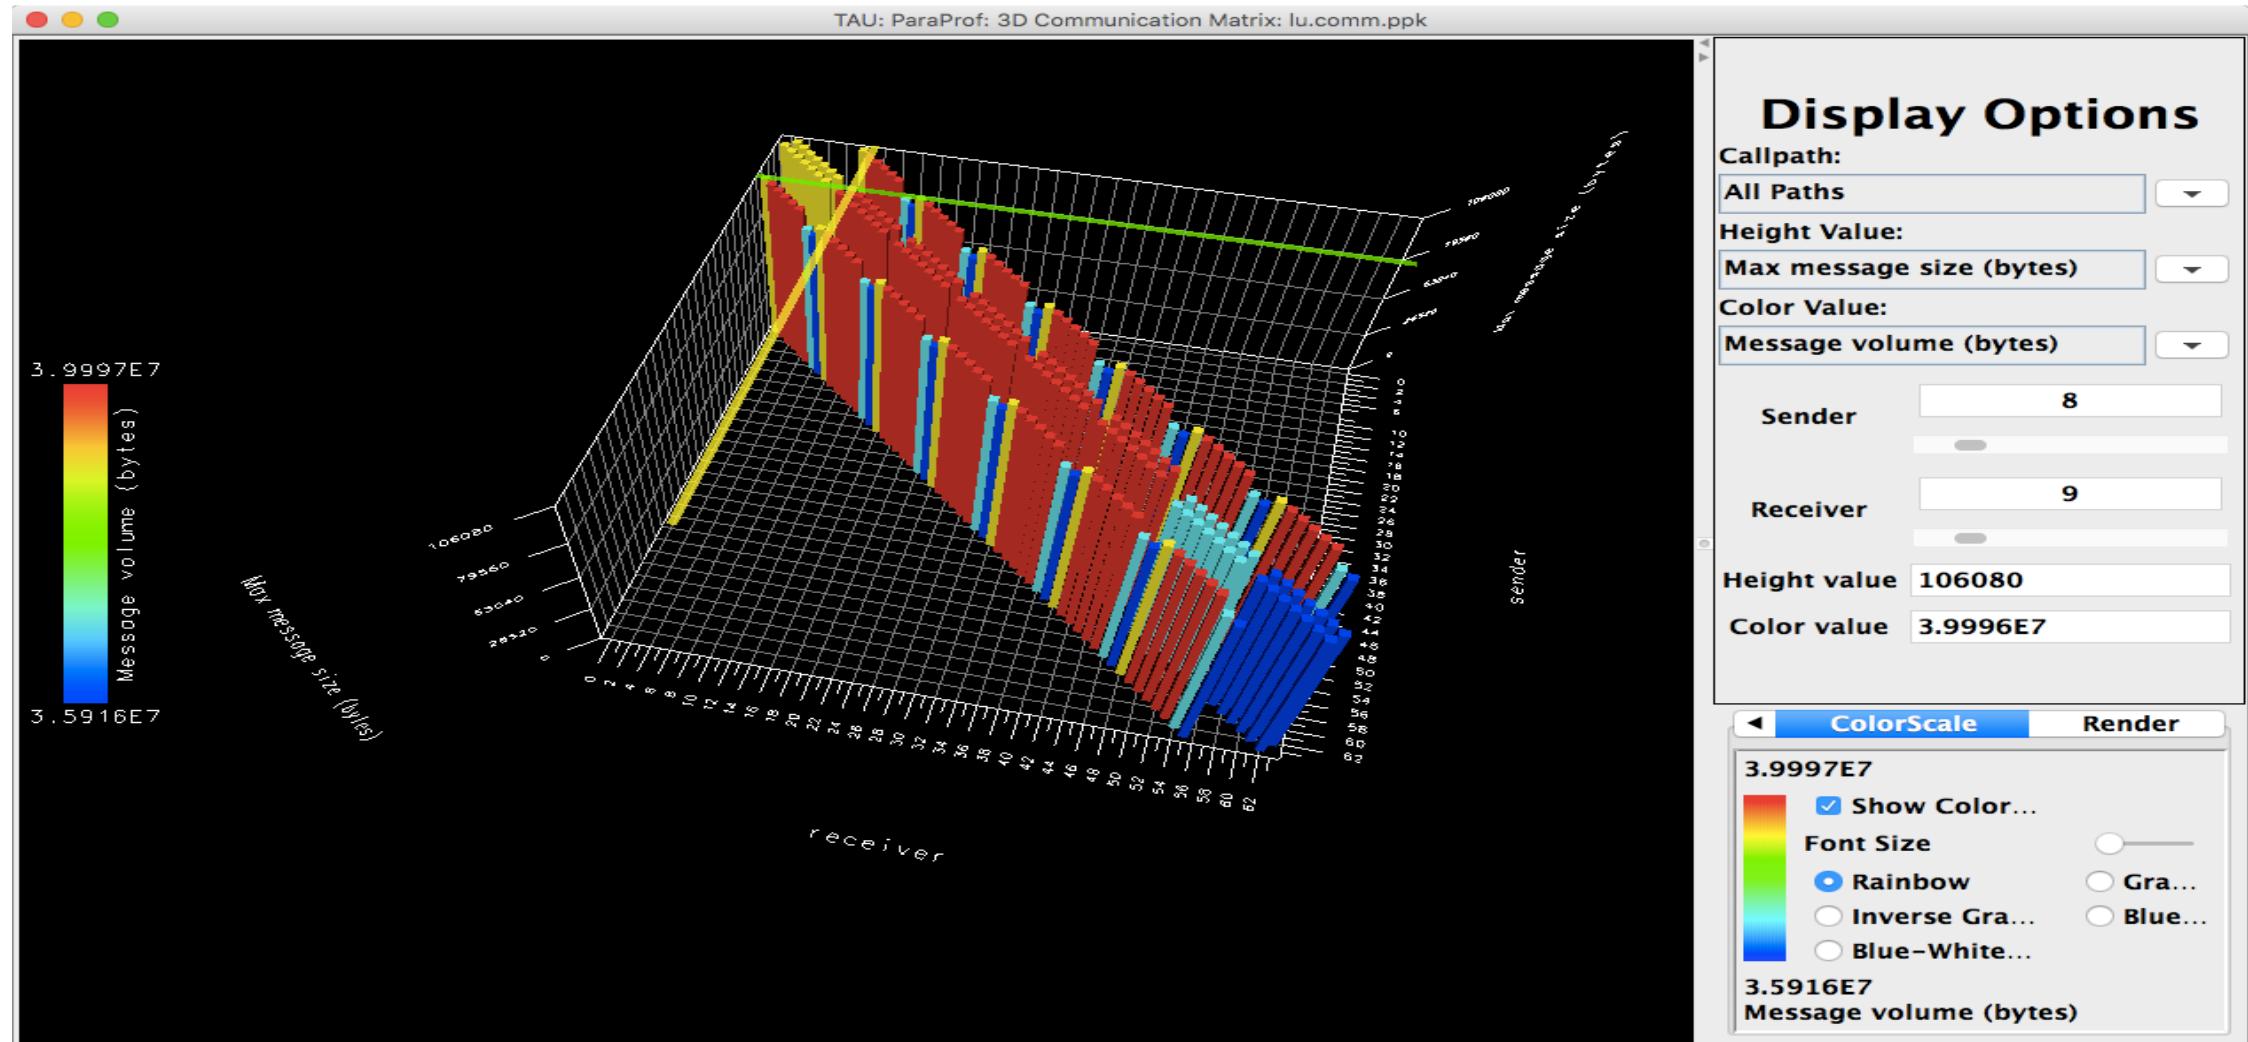


# TAU – ParaProf 3D Visualization



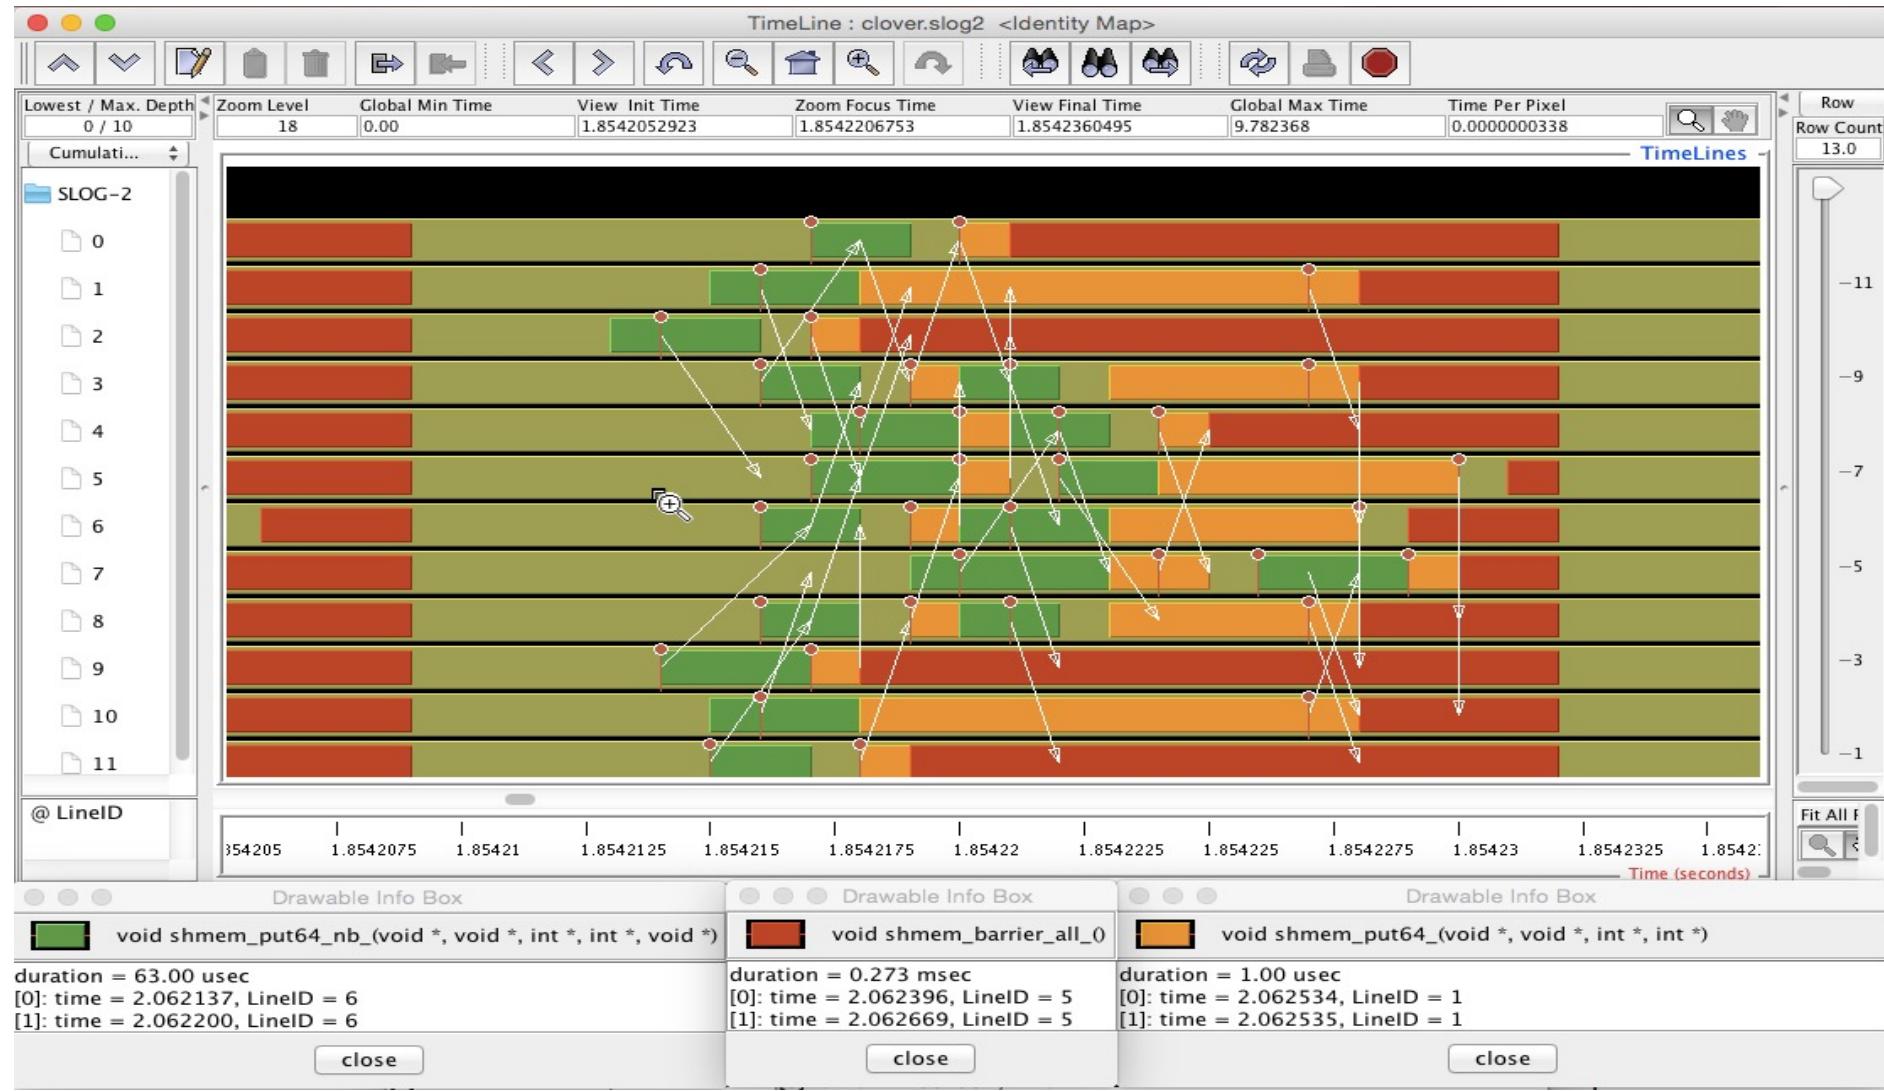
% paraprof app.ppk  
Windows -> 3D Visualization -> Bar Plot (right pane)

# TAU – 3D Communication Window

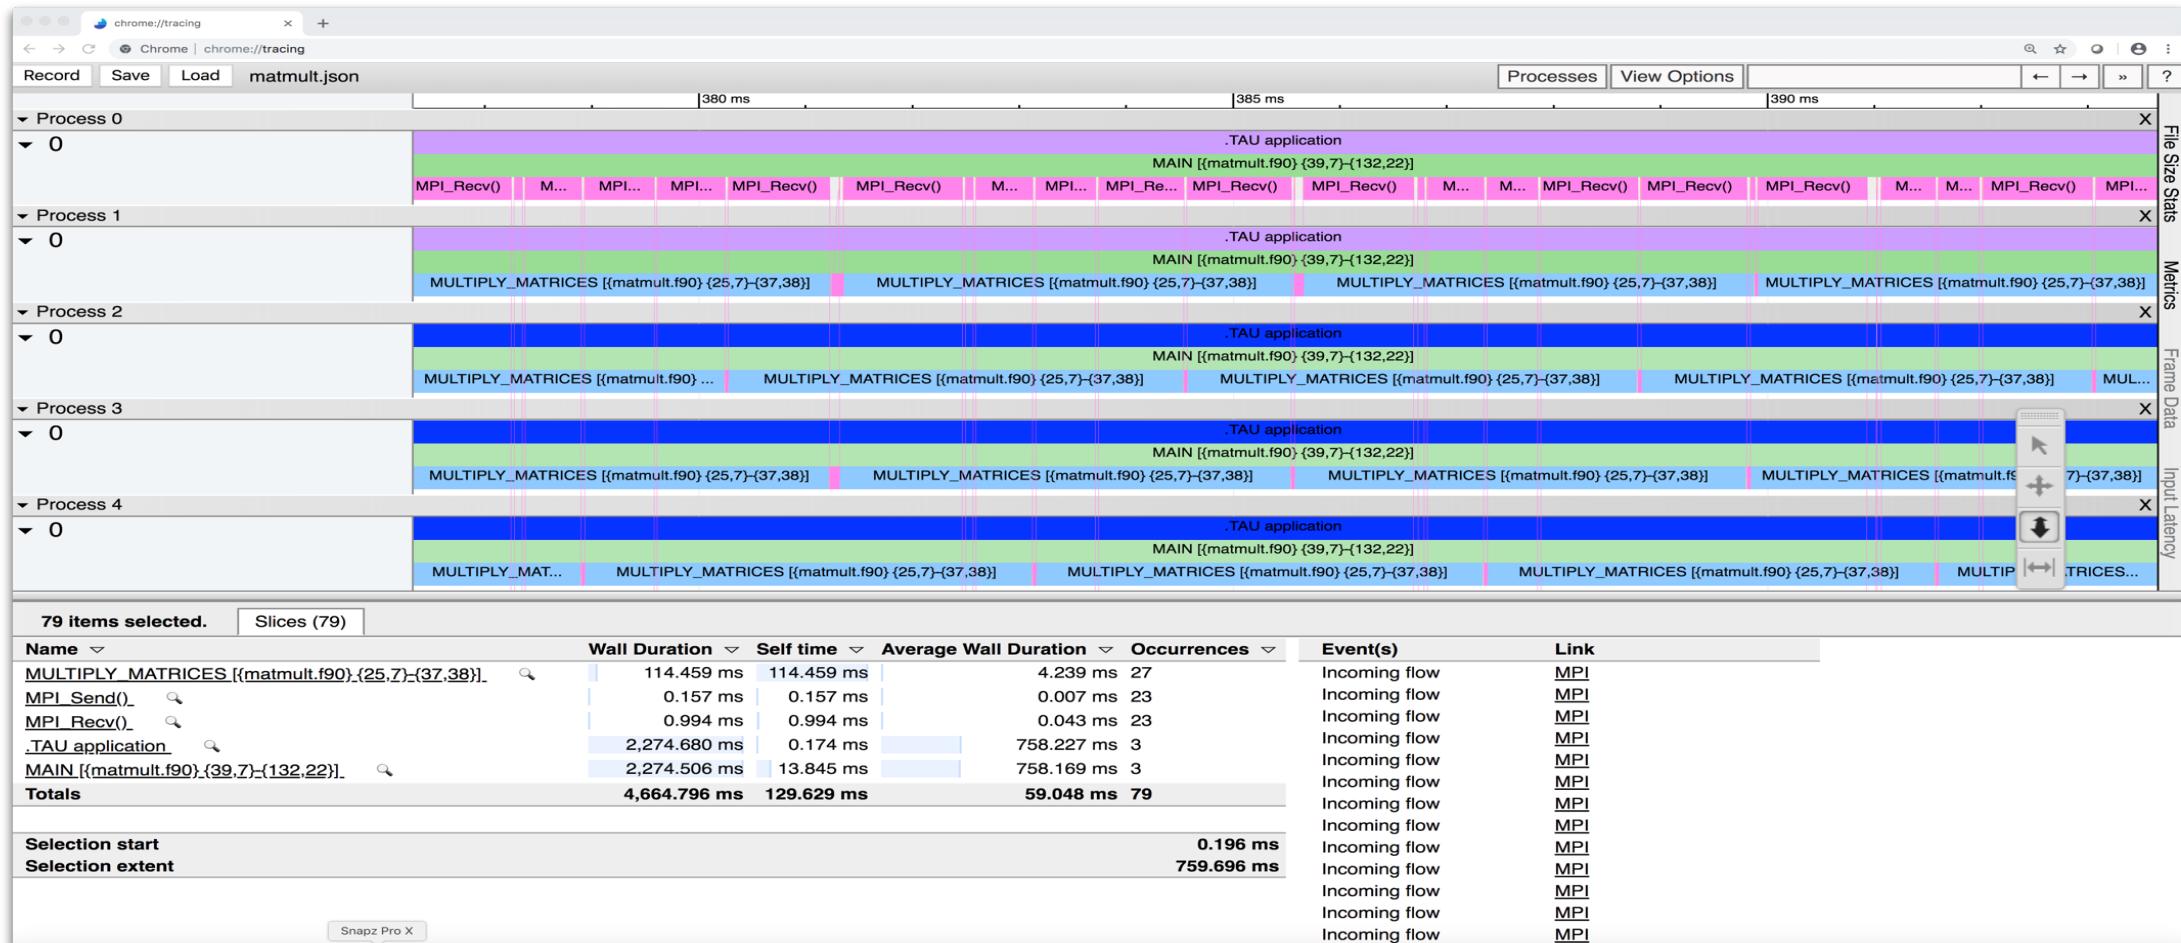


```
% export TAU_COMM_MATRIX=1; aprun ... tau_exec ./a.out  
% paraprof ; Windows -> 3D Communication Matrix
```

# Tracing: Jumpshot (ships with TAU)



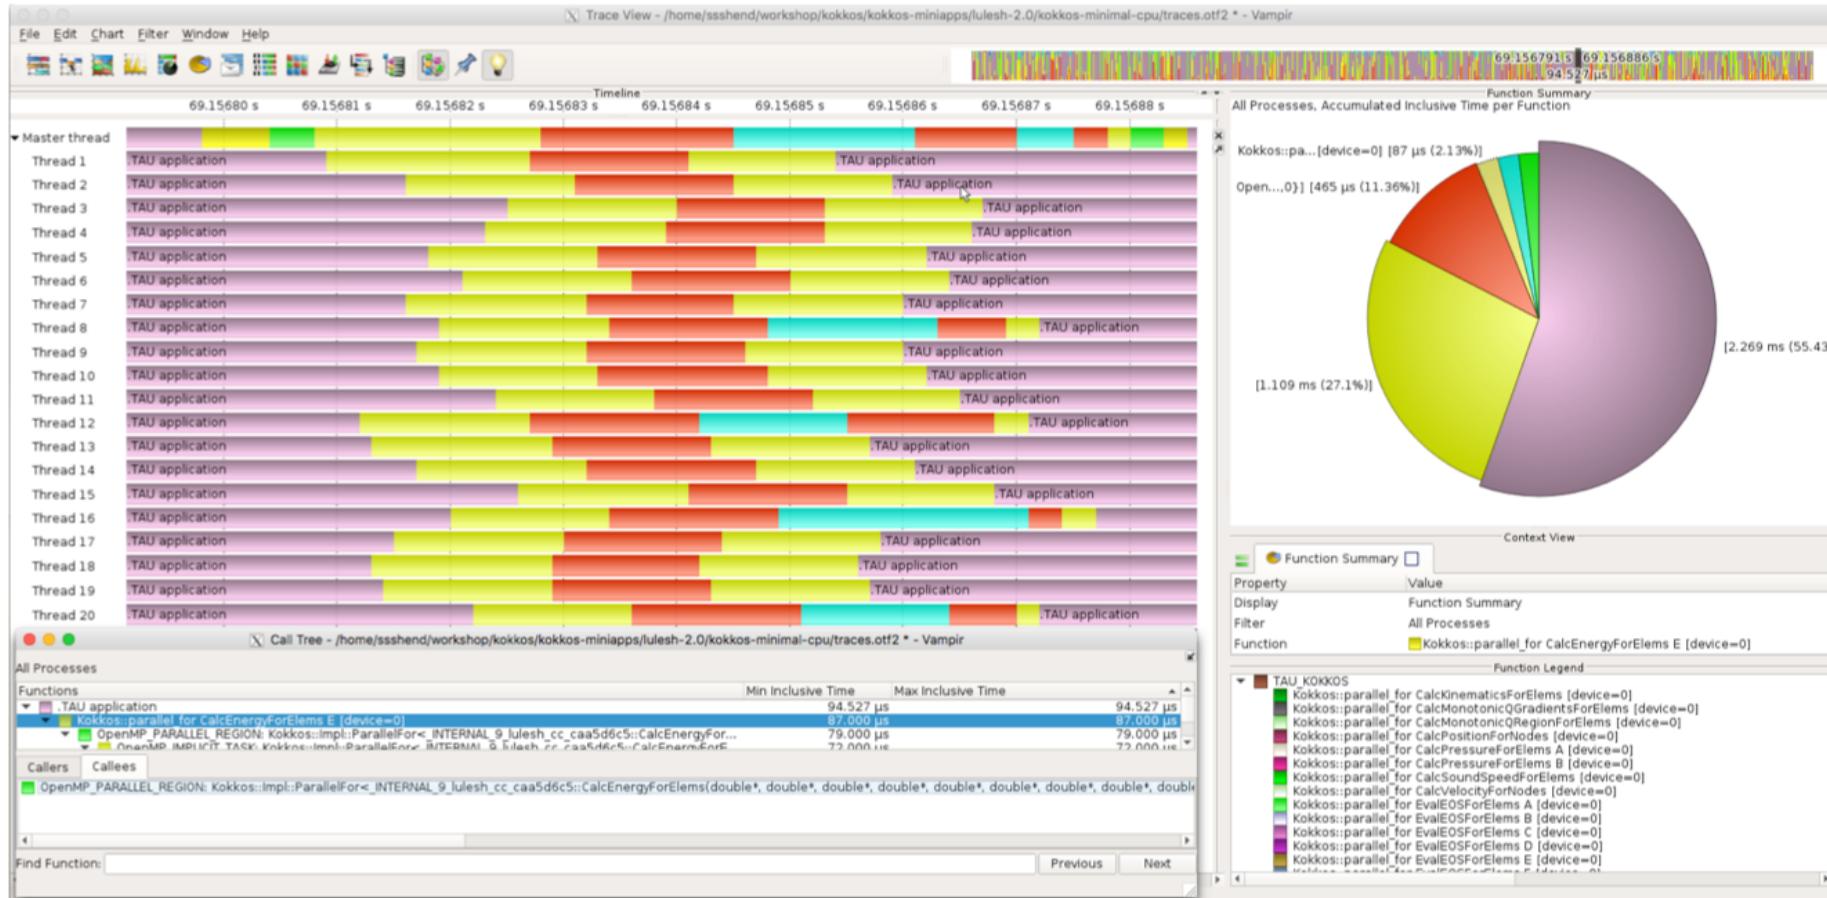
# Tracing: Chrome Browser



```
% export TAU_TRACE=1
% aprun -np 256 tau_exec ./a.out
% tau_treemerge.pl; tau_trace2json tau.trc tau.edf --chrome --ignoreatomic --o app.json
```

Chrome browser: chrome://tracing (Load -> app.json)

# Vampir [TU Dresden] Timeline: Kokkos



```
% export TAU_TRACE=1; export TAU_TRACE_FORMAT=otf2  
% tau_exec -omp ./a.out  
% vampir traces.otf2 &
```

# Kokkos

- Provides abstractions for node level parallelism (X in MPI+X)
- Productive, portable, and performant shared-memory programming model
- Helps you create single source performance portable codes
- Provides data abstractions
- C++ API for expressing parallelism in your program
- Aggressive compiler transformations using C++ templates
- Low level code targets backends such as OpenMP, Pthread, CUDA
- Creates a problem for performance evaluation tools
- Gap: performance data and higher-level abstractions
- Solution: Kokkos profiling API for mapping performance data

# Kokkos API use in ExaMiniMD

```
20. sameer@pegasus:~/pkgs/ORNL/DEMO/BUILD/ExaMiniMD-pthread/ExaMiniMD/src/comm_types (ssh)
void CommMPI::update_halo() {
    Kokkos::Profiling::pushRegion("Comm::update_halo"); ← pushRegion("Comm::update_halo")
    N_ghost = 0;
    s=*system;

    pack_buffer_update = t_buffer_update((T_X_FLOAT*)pack_buffer.data(),pack_indices_all.extent(1));
    unpack_buffer_update = t_buffer_update((T_X_FLOAT*)unpack_buffer.data(),pack_indices_all.extent(1));

    for(phase = 0; phase<6; phase++) {
        pack_indices = Kokkos::subview(pack_indices_all,phase,Kokkos::ALL());
        if(proc_grid[phase/2]>1) {

            Kokkos::parallel_for("CommMPI::halo_update_pack",
                Kokkos::RangePolicy<TagHaloUpdatePack, Kokkos::IndexType<T_INT> >(0,proc_num_send[phase]),
                *this);
            MPI_Request request;
            MPI_Status status;
            MPI_Irecv(unpack_buffer.data(),proc_num_recv[phase]*sizeof(T_X_FLOAT)*3/sizeof(int),MPI_INT, proc_neighbors_recv[phase],100002,MPI_COMM_WORLD,&request);
            MPI_Send (pack_buffer.data(),proc_num_send[phase]*sizeof(T_X_FLOAT)*3/sizeof(int),MPI_INT, proc_neighbors_send[phase],100002,MPI_COMM_WORLD);
            s = *system;
            MPI_Wait(&request,&status);
            const int count = proc_num_recv[phase];
            if(unpack_buffer_update.extent(0)<count) {
                unpack_buffer_update = t_buffer_update((T_X_FLOAT*)unpack_buffer.data(),count);
            }
            Kokkos::parallel_for("CommMPI::halo_update_unpack", ← Kokkos::parallel_for
                Kokkos::RangePolicy<TagHaloUpdateUnpack, Kokkos::IndexType<T_INT> >(0,proc_num_recv[phase]),
                *this);
        } else {
            //printf("HaloUpdateCopy: %i %i %i\n",phase,proc_num_send[phase],pack_indices.extent(0));
            Kokkos::parallel_for("CommMPI::halo_update_self",
                Kokkos::RangePolicy<TagHaloUpdateSelf, Kokkos::IndexType<T_INT> >(0,proc_num_send[phase]),
                *this);
        }
        N_ghost += proc_num_recv[phase]; ← popRegion
    }

    Kokkos::Profiling::popRegion();
};
```

# ExaMiniMD: TAU Phase

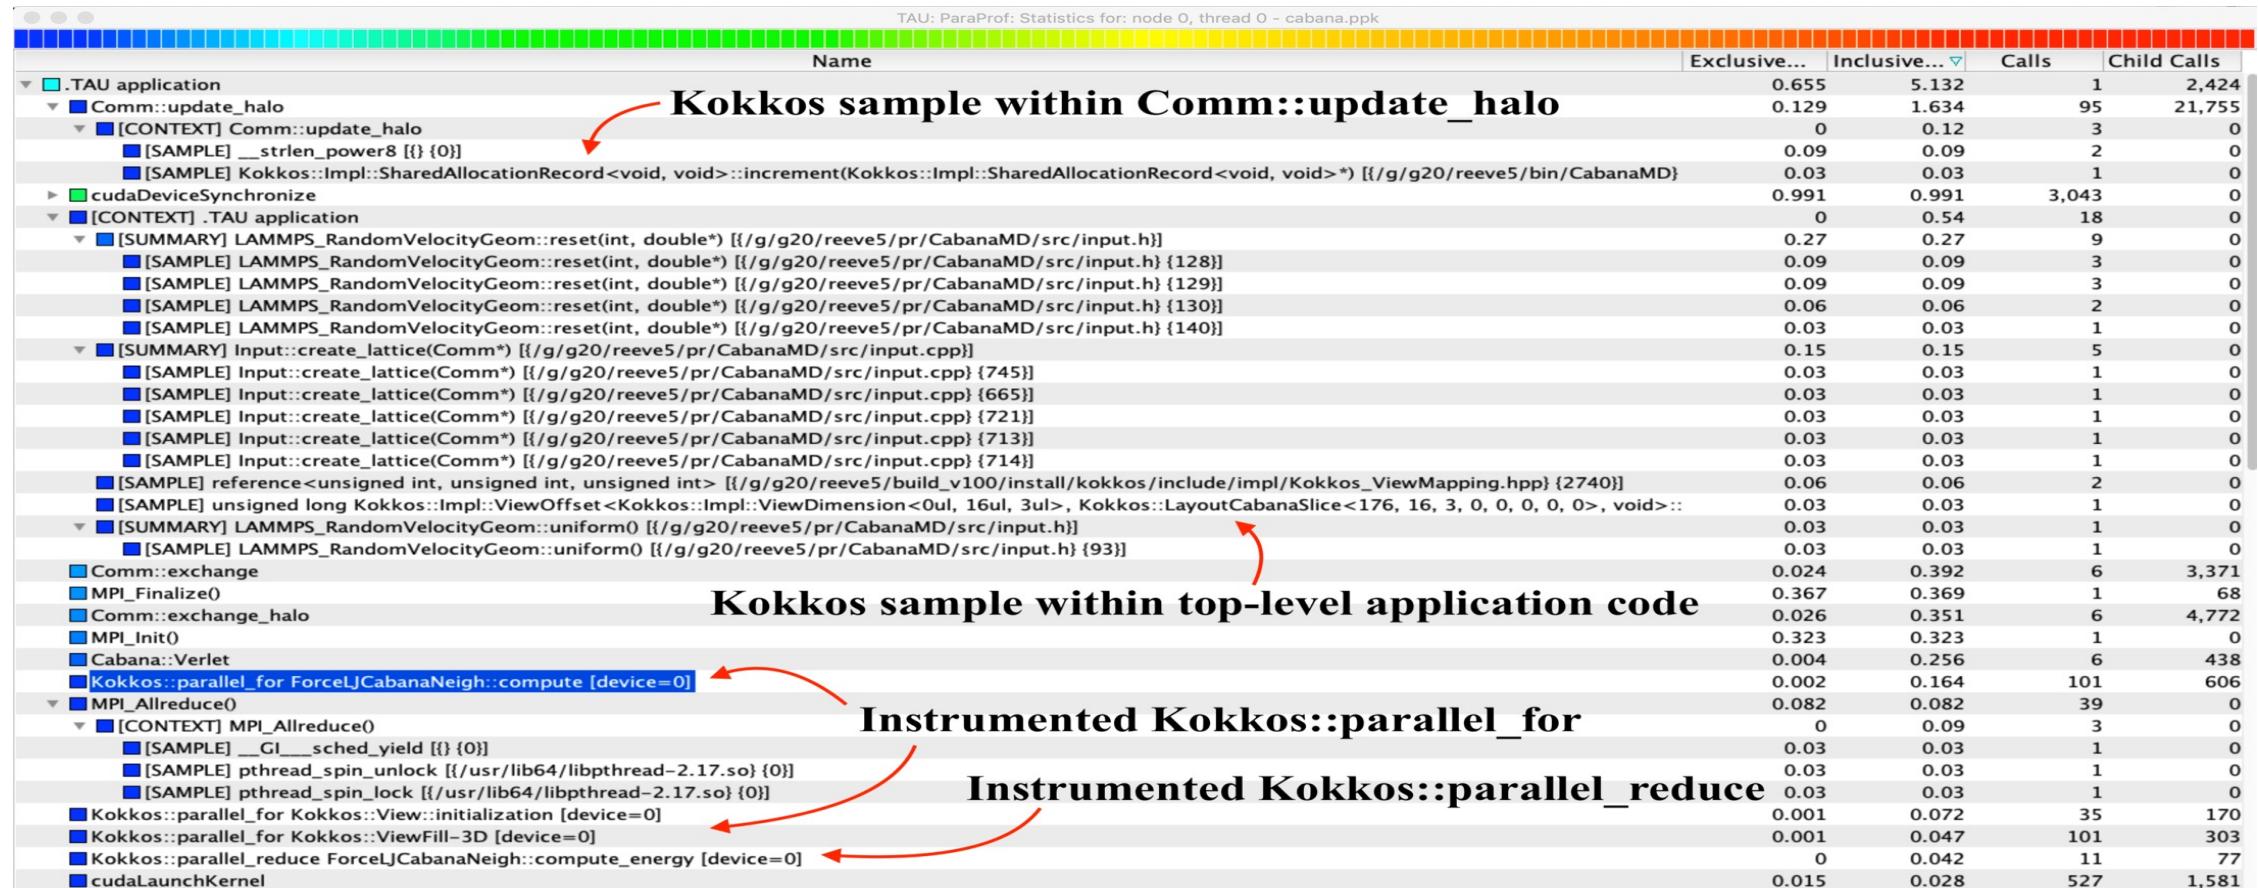
Name	Exclusive TIME	Inclusive TIME	Calls	Child Calls
.TAU application	0.143	96.743	1	832
Comm::exchange	0.001	0.967	6	142
Comm::exchange_halo	0.001	4.702	6	184
Comm::update_halo	0.004	31.347	95	1,330
Kokkos::parallel_for CommMPI::halo_update_pack [device=0]	0.002	0.506	190	190
Kokkos::parallel_for CommMPI::halo_update_self [device=0]	0.003	0.597	380	380
Kokkos::parallel_for CommMPI::halo_update_unpack [device=0]	0.002	0.97	190	190
MPI_Irecv()	0.001	0.001	190	0
MPI_Send()	29.268	29.268	190	0
MPI_Wait()	0.001	0.001	190	0
OpenMP_Implicit_Task	0.041	1.985	760	760
OpenMP_Parallel_Region parallel_for<Kokkos::RangePolicy<CommMPI::Ta	0	0.504	190	190
OpenMP_Parallel_Region parallel_for<Kokkos::RangePolicy<CommMPI::Ta	0.08	0.968	190	190
OpenMP_Parallel_Region void Kokkos::parallel_for<Kokkos::RangePolicy<	0.001	0.594	380	380
OpenMP_Sync_Region_Barrier parallel_for<Kokkos::RangePolicy<CommMF	0.489	0.489	190	0
OpenMP_Sync_Region_Barrier parallel_for<Kokkos::RangePolicy<CommMF	0.875	0.875	190	0
OpenMP_Sync_Region_Barrier void Kokkos::parallel_for<Kokkos::RangePol	0.58	0.58	380	0

Comm::update\_halo phase in TAU ParaProf's Thread Statistics Table

# ExaMiniMD: ParaProf Node Window

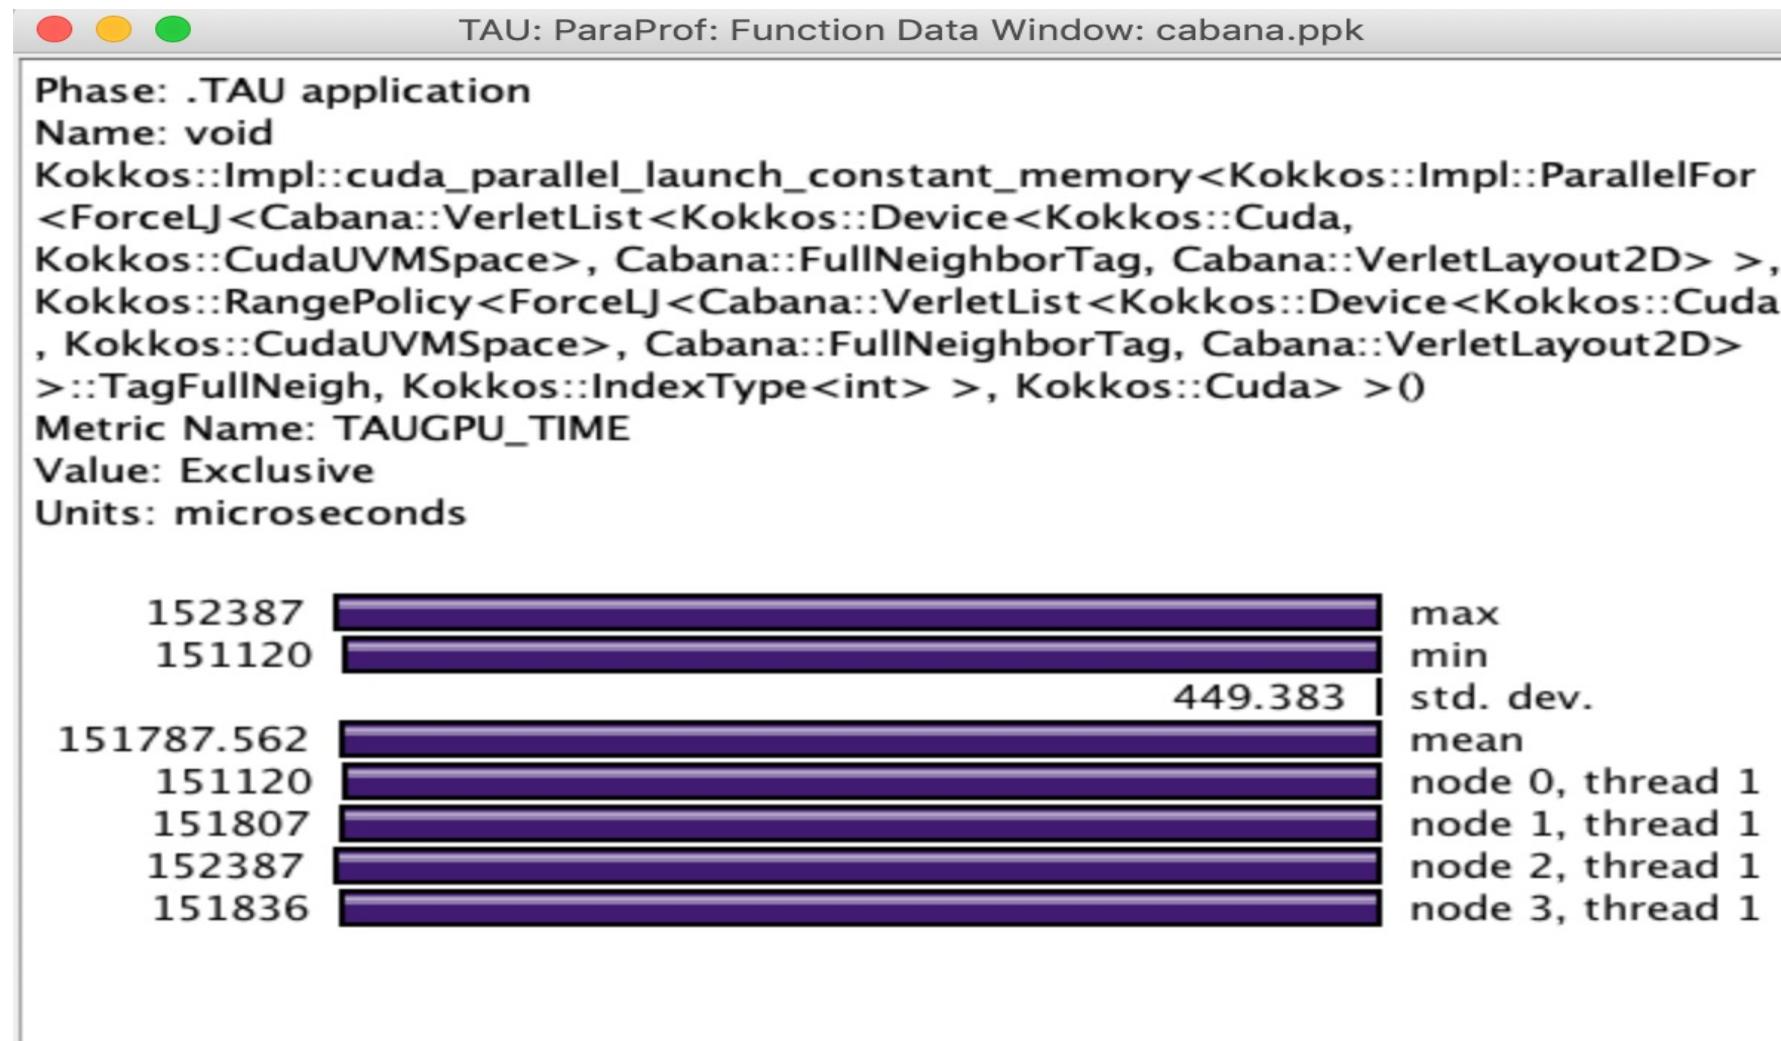


# Event-based Sampling (EBS): CabanaMD on an IBM AC922 with NVIDIA V100 GPUs



Event-based sampling (EBS) with Kokkos API

# CabanaMD: CUDA Events



# TAU's Support for Runtime Systems

## *MPI*

PMPI profiling interface

MPI\_T tools interface using performance and control variables

## *Pthread*

Captures time spent in routines per thread of execution

## *OpenMP*

OMPT tools interface to track salient OpenMP runtime events

Opari source rewriter

Preloading wrapper OpenMP runtime library when OMPT is not supported

## *OpenACC*

OpenACC instrumentation API

Track data transfers between host and device (per-variable)

Track time spent in kernels

# TAU's Support for Runtime Systems (contd.)

## *OpenCL*

- OpenCL profiling interface
- Track timings of kernels

## *CUDA*

- Cuda Profiling Tools Interface (CUPTI)
- Track data transfers between host and GPU
- Track access to uniform shared memory between host and GPU

## *ROCM*

- Rocprofiler and Roctracer instrumentation interfaces
- Track data transfers and kernel execution between host and GPU

## *Kokkos*

- Kokkos profiling API
- Push/pop interface for region, kernel execution interface

## *Python*

- Python interpreter instrumentation API
- Tracks Python routine transitions as well as Python to C transitions

# Examples of Multi-Level Instrumentation

## *MPI + OpenMP*

MPI\_T + PMPI + OMPT may be used to track MPI and OpenMP

## *MPI + CUDA*

PMPI + CUPTI interfaces

## *OpenCL + ROCm*

Rocprofiler + OpenCL instrumentation interfaces

## *Kokkos + OpenMP*

Kokkos profiling API + OMPT to transparently track events

## *Kokkos + pthread + MPI*

Kokkos + pthread wrapper interposition library + PMPI layer

## *Python + CUDA + MPI*

Python + CUPTI + pthread profiling interfaces (e.g., Tensorflow, PyTorch) + MPI

## *MPI + OpenCL*

PMPI + OpenCL profiling interfaces

# TAU Execution Command (tau\_exec)

Uninstrumented execution

```
% aprun -n 256 ./a.out
```

Track GPU operations

```
% aprun -np 256 tau_exec -rocm ./a.out  
% aprun -np 256 tau_exec -cupti ./a.out  
% aprun -np 256 tau_exec -opencl ./a.out  
% aprun -np 256 tau_exec -l0 ./a.out  
% aprun -np 256 tau_exec -openacc ./a.out
```

Track MPI performance

```
% aprun -n 256 tau_exec ./a.out
```

Track I/O, and MPI performance (MPI enabled by default)

```
% aprun -n 256 tau_exec -io ./a.out
```

Track OpenMP and MPI execution (using OMPT for Intel v19+ or Clang 8+)

```
% export TAU_OMPT_SUPPORT_LEVEL=full;  
% aprun -np 256 tau_exec -T ompt,v5,mpi -ompt ./a.out
```

Track memory operations

```
% export TAU_TRACK_MEMORY_LEAKS=1  
% aprun -np 256 tau_exec -memory_debug ./a.out (bounds check)
```

Use event based sampling (compile with -g)

```
% aprun -np 256 tau_exec -ebs ./a.out
```

 Also -ebs\_source=<PAPI\_COUNTER> -ebs\_period=<overflow\_count> -ebs\_resolution=<file | function | line>

# tau\_exec

```
$ tau_exec
```

```
Usage: tau_exec [options] [--] <exe> <exe options>
```

Options:

```
-v          Verbose mode
-s          Show what will be done but don't actually do anything (dryrun)
-qsub      Use qsub mode (BG/P only, see below)
-io         Track I/O
-memory    Track memory allocation/deallocation
-memory_debug Enable memory debugger
-cuda       Track GPU events via CUDA
-cupti      Track GPU events via CUPTI (Also see env. variable TAU_CUPTI_API)
-opencl     Track GPU events via OpenCL
-openacc    Track GPU events via OpenACC (currently PGI only)
-ompt       Track OpenMP events via OMPT interface
-armci      Track ARMCI events via PARMCI
-ebs        Enable event-based sampling
-ebs_period=<count> Sampling period (default 1000)
-ebs_source=<counter> Counter (default itimer)
-um         Enable Unified Memory events via CUPTI
-T <DISABLE,GNU,ICPC,MPI,OMPT,OPENMP,PAPI,PDT,PROFILE,PTHREAD,
    SCOREP,SERIAL>
    : Specify TAU tags
-loadlib=<file.so>   : Specify additional load library
-XrunTAUsh-<options> : Specify TAU library directly
-gdb        Run program in the gdb debugger
```

Notes:

Defaults if unspecified: -T MPI  
MPI is assumed unless SERIAL is specified

**tau\_exec**  
preloads the  
TAU wrapper  
libraries and  
performs  
measurements

No need to recompile  
the application!

# tau\_exec Example (continued)

Example:

```
aprun -np 2 tau_exec -T icpc,ompt,mpi -ompt ./a.out  
aprun -n 2 tau_exec -io ./a.out
```

Example - event-based sampling with samples taken every 1,000,000 FP instructions

```
aprun -n 8 tau_exec -ebs -ebs_period=1000000 -ebs_source=PAPI_FP_INS ./ring
```

Examples - GPU:

```
tau_exec -T serial,cupti -cupti ./matmult (Preferred for CUDA 4.1 or later)  
tau_exec -openacc ./a.out  
tau_exec -T serial -opencl ./a.out (OPENCL)  
aprunk -np 2 tau_exec -T mpi,cupti,papi -cupti -um ./a.out (Unified Virtual Memory in CUDA 6.0+)
```

qsub mode (IBM BG/Q only):

Original:

```
qsub -n 1 --mode smp -t 10 ./a.out
```

With TAU:

```
tau_exec -qsub -io -memory -- qsub -n 1 ... -t 10 ./a.out
```

Memory Debugging:

-memory option:

Tracks heap allocation/deallocation and memory leaks.

-memory\_debug option:

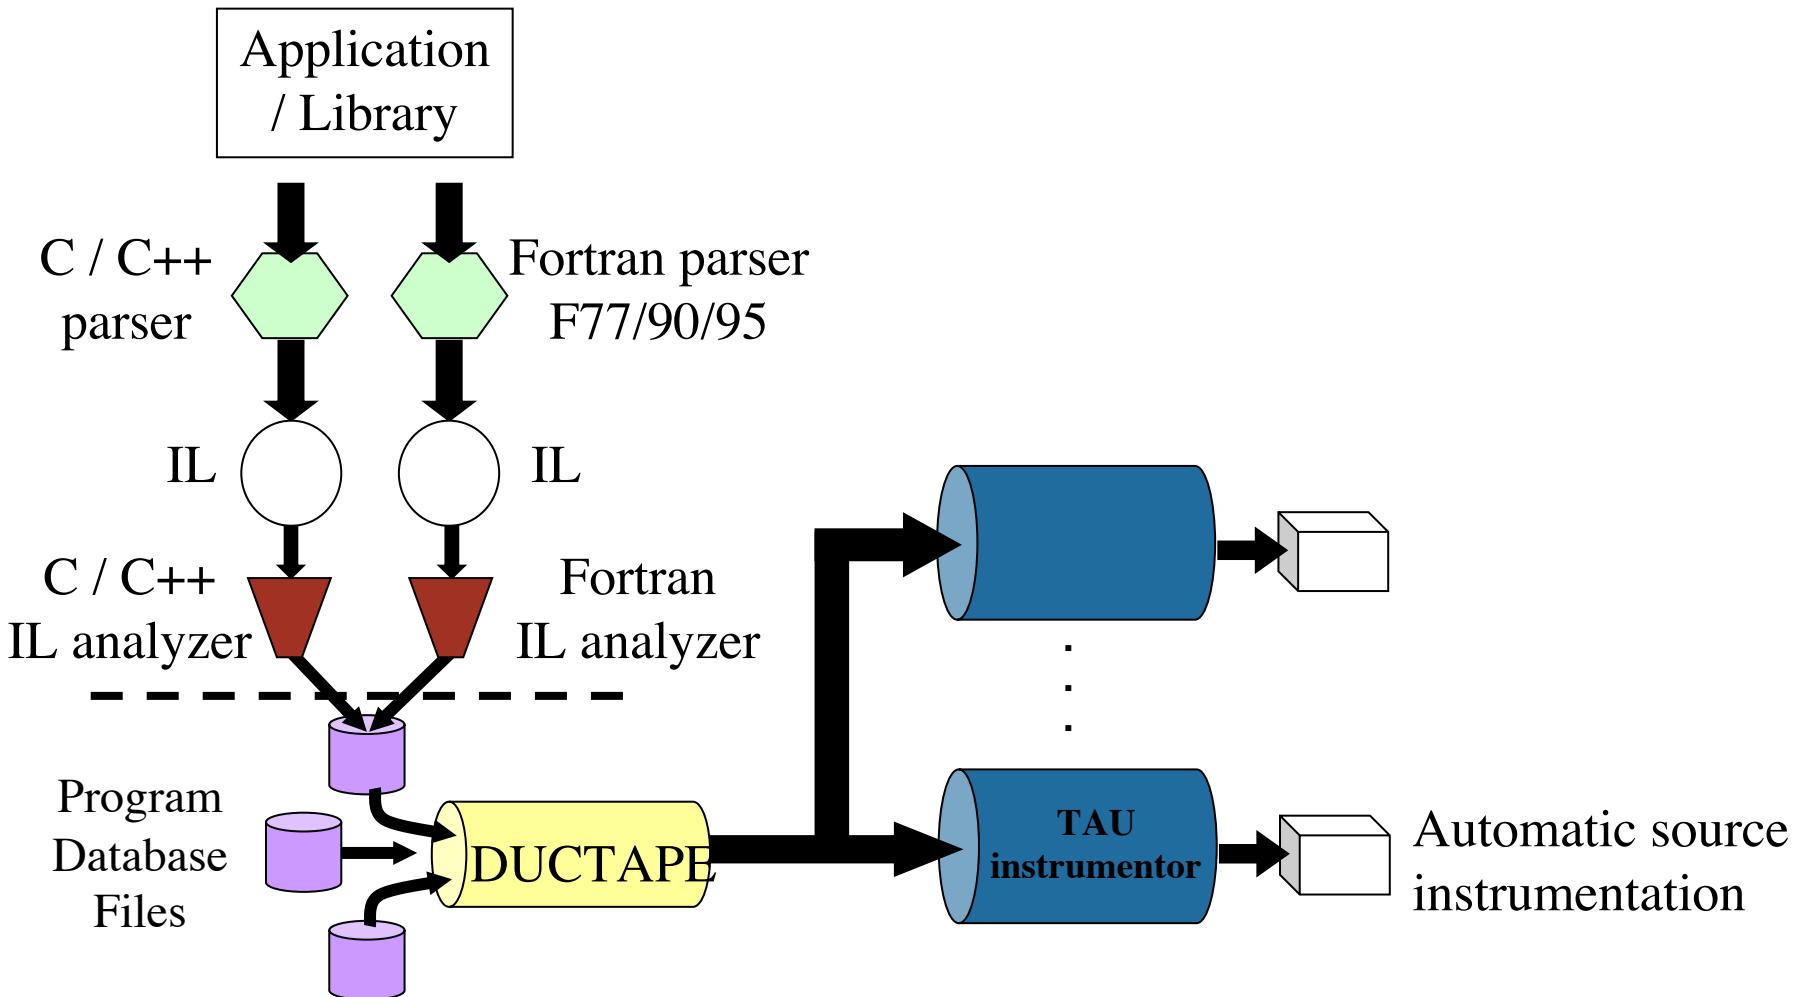
Detects memory leaks, checks for invalid alignment, and checks for array overflow. This is exactly like setting TAU\_TRACK\_MEMORY\_LEAKS=1 and TAU\_MEMDBG\_PROTECT\_ABOVE=1 and running with -memory

tau\_exec can enable event based sampling while launching the executable using environment var  
**TAU\_SAMPLING=1** or **tau\_exec -ebs**

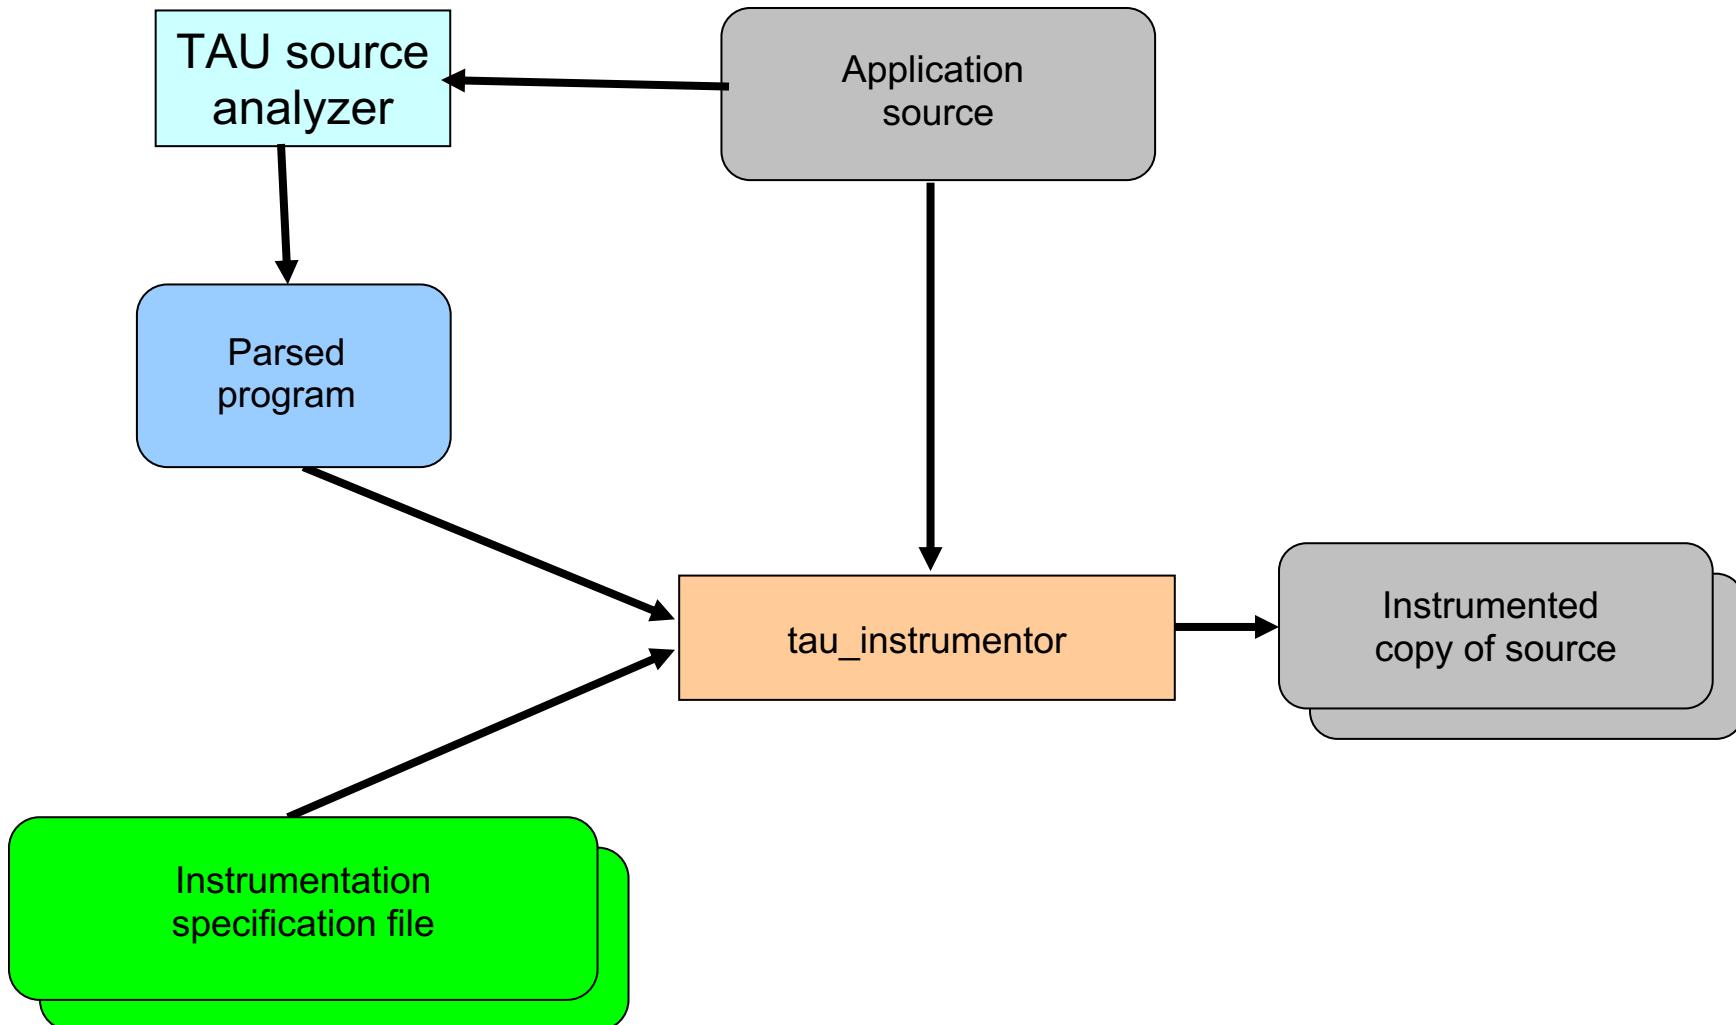
# Automatic Source Instrumentation in TAU using PDT



# TAU's Static Analysis System: Program Database Toolkit (PDT)



# PDT: automatic source instrumentation



# Installing TAU

- Installing PDT:
  - wget [http://tau.uoregon.edu/pdt\\_lite.tgz](http://tau.uoregon.edu/pdt_lite.tgz)
  - ./configure –prefix=<dir>; make ; make install
- Installing TAU:
  - wget <http://tau.uoregon.edu/tau.tgz>
  - ./configure -mpi –pdt=<dir> -bfd=download –unwind=download –iowrapper ...
  - make install
- Using TAU:
  - export TAU\_MAKEFILE=<taudir>/ibm64linux/lib/Makefile.tau-<TAGS>
  - make CC=tau\_cc.sh CXX=tau\_cxx.sh F90=tau\_f90.sh

# Installing TAU on Laptops

- Installing TAU under Mac OS X:
  - wget <http://tau.uoregon.edu/tau.dmg>
  - Install tau.dmg
- Installing TAU under Linux
  - <http://tau.uoregon.edu/tau.exe>
- Installing TAU under Linux
  - <http://tau.uoregon.edu/tau.tgz>
  - ./configure; make install
  - export PATH=<taudir>/x86\_64/bin:\$PATH

# Source Instrumentation in TAU

- TAU supports several compilers, measurement, and thread options
  - Intel, GNU, Clang, PGI compilers, profiling with hardware counters using PAPI, MPI library, OpenMP...
  - Each measurement configuration corresponds to a unique stub makefile (configuration file) and library that is generated when you configure it
- To instrument source code automatically using PDT
  - Choose an appropriate TAU stub makefile in <arch>/lib:
    - % module load tau
    - % export TAU\_MAKEFILE=<tau\_root>/craycnl/lib/Makefile.tau-<options>
    - % export TAU\_OPTIONS=' -optVerbose ...' (see tau\_compiler.sh )
  - Use tau\_f90.sh, tau\_cxx.sh, tau\_upc.sh, tau\_caf.sh, or tau\_cc.sh as F90, C++, UPC, CAF, or C compilers respectively:
    - % ftn foo.f90 changes to
    - % **tau\_f90.sh** foo.f90
  - Set runtime environment variables, execute application and analyze performance data:
    - % pprof (for text based profile display)
    - % paraprof (for GUI)

# Different Makefiles for TAU Compiler on Summit at ORNL

```
% module load tau
% ls $TAU_DIR/lib/Makefile*
/sw/summit/tau/2.29.1/ibm64linux/lib/Makefile.tau-xl_16.1.1-5-papi-cupti
/sw/summit/tau/2.29.1/ibm64linux/lib/Makefile.tau-xl_16.1.1-5-papi-cupti-openmp-opari
/sw/summit/tau/2.29.1/ibm64linux/lib/Makefile.tau-xl_16.1.1-5-papi-mpi
/sw/summit/tau/2.29.1/ibm64linux/lib/Makefile.tau-xl_16.1.1-5-papi-mpi-cupti
/sw/summit/tau/2.29.1/ibm64linux/lib/Makefile.tau-xl_16.1.1-5-papi-mpi-cupti-openmp-opari
/sw/summit/tau/2.29.1/ibm64linux/lib/Makefile.tau-xl_16.1.1-5-papi-mpi-openmp-opari
/sw/summit/tau/2.29.1/ibm64linux/lib/Makefile.tau-xl_16.1.1-6-papi-cupti
/sw/summit/tau/2.29.1/ibm64linux/lib/Makefile.tau-xl_16.1.1-6-papi-cupti-openmp-opari
/sw/summit/tau/2.29.1/ibm64linux/lib/Makefile.tau-xl_16.1.1-6-papi-mpi
/sw/summit/tau/2.29.1/ibm64linux/lib/Makefile.tau-xl_16.1.1-6-papi-mpi-cupti
/sw/summit/tau/2.29.1/ibm64linux/lib/Makefile.tau-xl_16.1.1-6-papi-mpi-cupti-openmp-opari
/sw/summit/tau/2.29.1/ibm64linux/lib/Makefile.tau-xl_16.1.1-6-papi-mpi-openmp-opari
/sw/summit/tau/2.29.1/ibm64linux/lib/Makefile.tau-gcc_8.1.1-papi-gnu-cupti
/sw/summit/tau/2.29.1/ibm64linux/lib/Makefile.tau-gcc_8.1.1-papi-gnu-cupti-pdt
/sw/summit/tau/2.29.1/ibm64linux/lib/Makefile.tau-gcc_8.1.1-papi-gnu-mpi
/sw/summit/tau/2.29.1/ibm64linux/lib/Makefile.tau-gcc_8.1.1-papi-gnu-mpi-cupti
...
...
```

# Different Makefiles for TAU Compiler on Theta at ALCF

```
% module load tau
% ls $TAU_DIR/lib/Makefile*
/soft/perftools/tau/tau-2.30.1/craycnl/lib/Makefile.tau-intel-papi-mpi-pdt
/soft/perftools/tau/tau-2.30.1/craycnl/lib/Makefile.tau-intel-papi-mpi-pthread-pdt
/soft/perftools/tau/tau-2.30.1/craycnl/lib/Makefile.tau-intel-papi-ompt-v5-mpi-pdt-openmp
/soft/perftools/tau/tau-2.30.1/craycnl/lib/Makefile.tau-intel-papi-ompt-v5-pdt-openmp
/soft/perftools/tau/tau-2.30.1/craycnl/lib/Makefile.tau-intel-papi-pthread-pdt
...
...
```

# Configuration tags for tau\_exec

```
% ./configure -pdt=<dir> -mpi -papi=<dir>; make install
```

Creates in <taudir>/<arch>/lib:

```
Makefile.tau-papi-mpi-pdt  
shared-papi-mpi-pdt/libTAU.so
```

```
% ./configure -pdt=<dir> -mpi; make install creates
```

```
Makefile.tau-mpi-pdt  
shared-mpi-pdt/libTAU.so
```

To explicitly choose preloading of shared-<options>/libTAU.so change:

```
% aprun -np 8 ./a.out      to
```

```
% aprun -np 8 tau_exec -T <comma_separated_options> ./a.out
```

```
% aprun -np 8 tau_exec -T papi,mpi,pdt ./a.out
```

Preloads <taudir>/<arch>/shared-papi-mpi-pdt/libTAU.so

```
% aprun -np 8 tau_exec -T papi ./a.out
```

Preloads <taudir>/<arch>/shared-papi-mpi-pdt/libTAU.so by matching.

```
% aprun -np 8 tau_exec -T papi,mpi,pdt -s ./a.out
```

Does not execute the program. Just displays the library that it will preload if executed without the -s option.

NOTE: -mpi configuration is selected by default. Use -T serial for Sequential programs.

# Compile-Time Options

Optional parameters for the TAU\_OPTIONS environment variable:

% tau\_compiler.sh

-optVerbose	Turn on verbose debugging messages
-optComplInst	Use compiler based instrumentation
-optNoComplInst	Do not revert to compiler instrumentation if source instrumentation fails.
-optTrackIO	Wrap POSIX I/O call and calculates vol/bw of I/O operations (configure TAU with <i>-iowrapper</i> )
-optTrackGOMP	Enable tracking GNU OpenMP runtime layer (used without <i>-opari</i> )
-optMemDbg	Enable runtime bounds checking (see TAU_MEMDBG_* env vars)
-optKeepFiles	Does not remove intermediate .pdb and .inst.* files
-optPreProcess	Preprocess sources (OpenMP, Fortran) before instrumentation
-optTauSelectFile=<file>	Specify selective instrumentation file for <i>tau_instrumentor</i>
-optTauWrapFile=<file>	Specify path to <i>link_options.tau</i> generated by <i>tau_gen_wrapper</i>
-optHeaderInst	Enable Instrumentation of headers
-optTrackUPCR	Track UPC runtime layer routines (used with <i>tau_upc.sh</i> )
-optLinking=""	Options passed to the linker. Typically \$(TAU_MPI_FLIBS) \$(TAU_LIBS) \$(TAU_CXXLIBS)
-optCompile=""	Options passed to the compiler. Typically \$(TAU_MPI_INCLUDE) \$(TAU_INCLUDE) \$(TAU_DEFS)
-optPdtF95Opts=""	Add options for Fortran parser in PDT (f95parse/gfparse) ...

# Compile-Time Options (contd.)

Optional parameters for the TAU\_OPTIONS environment variable:

% tau\_compiler.sh

-optShared	Use TAU's shared library (libTAU.so) instead of static library (default)
-optPdtCxxOpts=""	Options for C++ parser in PDT (cxxparse).
-optPdtF90Parser=""	Specify a different Fortran parser
-optPdtCleanscapeParser	Specify the Cleanscape Fortran parser instead of GNU gparser
-optTau=""	Specify options to the tau_instrumentor
-optTrackDMAPP	Enable instrumentation of low-level DMAPP API calls on Cray
-optTrackPthread	Enable instrumentation of pthread calls

See tau\_compiler.sh for a full list of TAU\_OPTIONS.

...

# Using TAU\_OPTIONS

To use the compiler based instrumentation instead of PDT (source-based):

```
% export TAU_OPTIONS= '-optComInst -optVerbose'
```

If your Fortran code uses C preprocessor directives (#include, #ifdef, #endif):

```
% export TAU_OPTIONS= '-optPreProcess -optVerbose -optDetectMemoryLeaks'
```

To use an instrumentation specification file:

```
% export TAU_OPTIONS= '-optTauSelectFile=select.tau -optVerbose -optPreProcess'  
% cat select.tau  
BEGIN_INSTRUMENT_SECTION  
loops routine="#"  
# this statement instruments all outer loops in all routines. # is wildcard as well as comment in first column.  
END_INSTRUMENT_SECTION
```

# Selective Instrumentation File With Program Database Toolkit (PDT)

To use an instrumentation specification file for source instrumentation:

```
% export TAU_OPTIONS=' -optTauSelectFile=/path/to/select.tau --optVerbose '
```

```
% cat select.tau
```

```
BEGIN_EXCLUDE_LIST
```

```
BINVCRHS
```

```
MATMUL_SUB
```

```
MATVEC_SUB
```

```
EXACT_SOLUTION
```

```
BINVRHS
```

```
LHS#INIT
```

```
TIMER_#
```

```
END_EXCLUDE_LIST
```

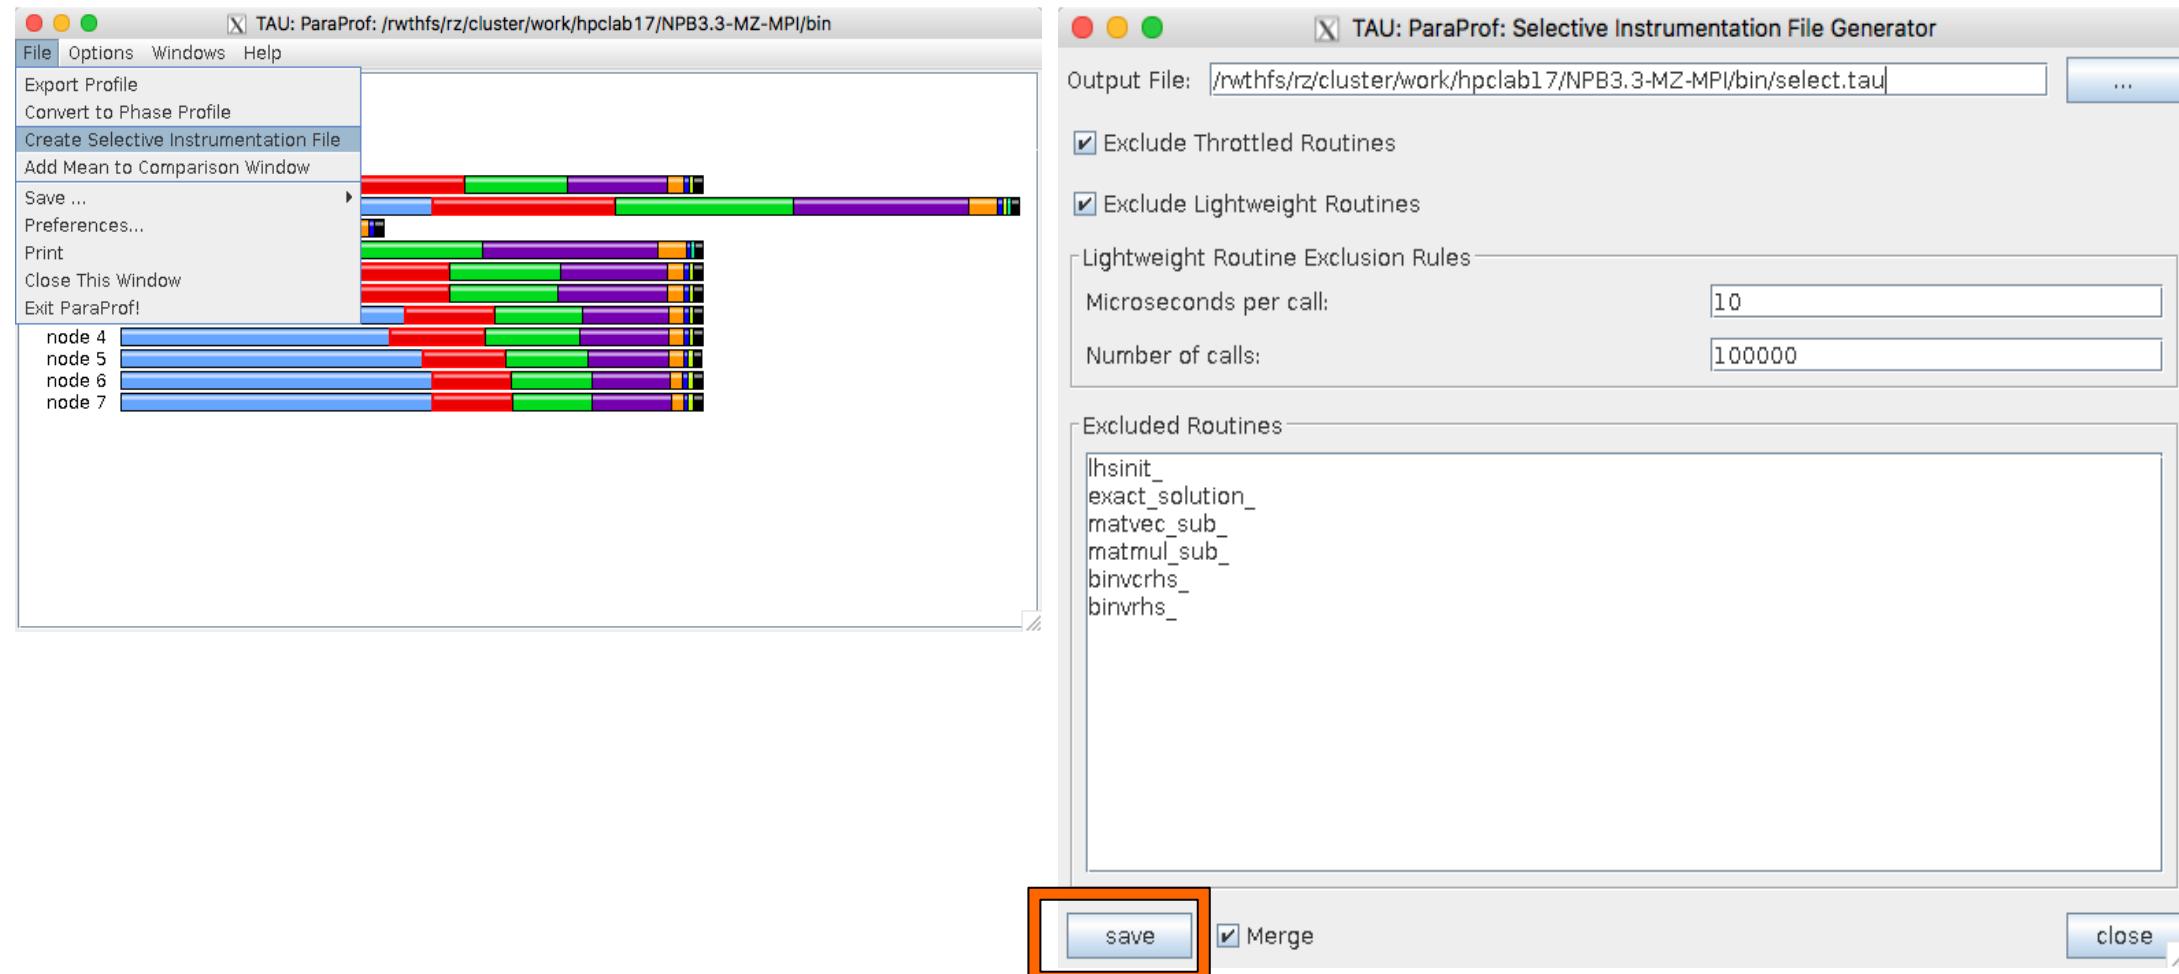
**NOTE:** paraprof can create this file from an earlier execution for you.

File -> Create Selective Instrumentation File -> save

Selective instrumentation at runtime:

```
% export TAU_SELECT_FILE=select.tau
```

# Create a Selective Instrumentation File, Re-instrument, Re-run



# TAU's Runtime Environment Variables

Environment Variable	Default	Description
TAU_TRACE	0	Setting to 1 turns on tracing
TAU_CALLPATH	0	Setting to 1 turns on callpath profiling
TAU_TRACK_MEMORY_FOOTPRINT	0	Setting to 1 turns on tracking memory usage by sampling periodically the resident set size and high water mark of memory usage
TAU_TRACK_POWER	0	Tracks power usage by sampling periodically.
TAU_CALLPATH_DEPTH	2	Specifies depth of callpath. Setting to 0 generates no callpath or routine information, setting to 1 generates flat profile and context events have just parent information (e.g., Heap Entry: foo)
TAU_SAMPLING	1	Setting to 1 enables event-based sampling.
TAU_TRACK_SIGNALS	0	Setting to 1 generate debugging callstack info when a program crashes
TAU_COMM_MATRIX	0	Setting to 1 generates communication matrix display using context events
TAU_THROTTLE	1	Setting to 0 turns off throttling. Throttles instrumentation in lightweight routines that are called frequently
TAU_THROTTLE_NUMCALLS	100000	Specifies the number of calls before testing for throttling
TAU_THROTTLE_PERCALL	10	Specifies value in microseconds. Throttle a routine if it is called over 100000 times and takes less than 10 usec of inclusive time per call
TAU_CALLSITE	0	Setting to 1 enables callsite profiling that shows where an instrumented function was called. Also compatible with tracing.
TAU_PROFILE_FORMAT	Profile	Setting to "merged" generates a single file. "snapshot" generates xml format
TAU_METRICS	TIME	Setting to a comma separated list generates other metrics. (e.g., ENERGY,TIME,P_VIRTUAL_TIME,PAPI_FP_INS,PAPI_NATIVE_<event>:<subevent>)

# Runtime Environment Variables

Environment Variable	Default	Description
TAU_TRACE	0	Setting to 1 turns on tracing
TAU_TRACE_FORMAT	Default	Setting to “otf2” turns on TAU’s native OTF2 trace generation (configure with –otf=download)
TAU_EBS_UNWIND	0	Setting to 1 turns on unwinding the callstack during sampling (use with tau_exec –ebs or TAU_SAMPLING=1)
TAU_EBS_RESOLUTION	line	Setting to “function” or “file” changes the sampling resolution to function or file level respectively.
TAU_TRACK_LOAD	0	Setting to 1 tracks system load on the node
TAU_SELECT_FILE	Default	Setting to a file name, enables selective instrumentation based on exclude/include lists specified in the file.
TAU_OMPT_SUPPORT_LEVEL	basic	Setting to “full” improves resolution of OMPT TR6 regions on threads 1.. N-1. Also, “lowoverhead” option is available.
TAU_OMPT_RESOLVE_ADDRESS_EAGERLY	1	Setting to 1 is necessary for event based sampling to resolve addresses with OMPT. Setting to 0 allows the user to do offline address translation.

# Runtime Environment Variables

Environment Variable	Default	Description
TAU_TRACK_MEMORY_LEAKS	0	Tracks allocates that were not de-allocated (needs –optMemDbg or tau_exec –memory)
TAU_EBS_SOURCE	TIME	Allows using PAPI hardware counters for periodic interrupts for EBS (e.g., TAU_EBS_SOURCE=PAPI_TOT_INS when TAU_SAMPLING=1)
TAU_EBS_PERIOD	100000	Specifies the overflow count for interrupts
TAU_MEMDBG_ALLOC_MIN/MAX	0	Byte size minimum and maximum subject to bounds checking (used with TAU_MEMDBG_PROTECT_*)
TAU_MEMDBG_OVERHEAD	0	Specifies the number of bytes for TAU's memory overhead for memory debugging.
TAU_MEMDBG_PROTECT_BELOW/ABOVE	0	Setting to 1 enables tracking runtime bounds checking below or above the array bounds (requires –optMemDbg while building or tau_exec –memory)
TAU_MEMDBG_ZERO_MALLOC	0	Setting to 1 enables tracking zero byte allocations as invalid memory allocations.
TAU_MEMDBG_PROTECT_FREE	0	Setting to 1 detects invalid accesses to deallocated memory that should not be referenced until it is reallocated (requires –optMemDbg or tau_exec –memory)
TAU_MEMDBG_ATTEMPT_CONTINUE	0	Setting to 1 allows TAU to record and continue execution when a memory error occurs at runtime.
TAU_MEMDBG_FILL_GAP	Undefined	Initial value for gap bytes
TAU_MEMDBG_ALIGNMENT	Sizeof(int)	Byte alignment for memory allocations
TAU_EVENT_THRESHOLD	0.5	Define a threshold value (e.g., .25 is 25%) to trigger marker events for min/max

# Extreme-scale Scientific Software Stack (E4S)

<https://e4s.io>



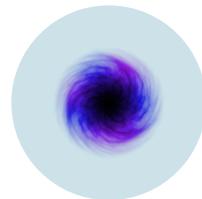
# E4S: Better Quality, Documentation, Test, Integration, Delivery, Build & Use

*Delivering HPC software to facilities, vendors, agencies, industry, international partners in a brand-new way*



## Community Policies

Commitment to software quality



## DocPortal

Single portal to all E4S product info



## Portfolio testing

Especially leadership platforms



## Curated collection

The end of dependency hell



## Quarterly releases

Release 1.2 – November



## Build caches

10X build time improvement



## Turnkey stack

A new user experience



<https://e4s.io>



## E4S Strategy Group

US agencies, industry, international

# E4S: Extreme-scale Scientific Software Stack

- Curated, Spack based software distribution
- Spack binary build caches for bare-metal installs
  - x86\_64, ppc64le (IBM Power 9), and aarch64 (ARM64)
- Container images on DockerHub and E4S website of pre-built binaries of ECP ST products
- Base images and full featured containers (with GPU support)
- GitHub recipes for creating custom images from base images
- GitLab integration for building E4S images
- E4S validation test suite on GitHub
- E4S-cl container launcher tool for MPI substitution in applications using MPICH ABI
- E4S VirtualBox image with support for container runtimes
  - Docker
  - Singularity
  - Shifter
  - Charliecloud
- AWS and GCP images to deploy E4S

<https://e4s.io>

# Spack is a flexible package manager for HPC

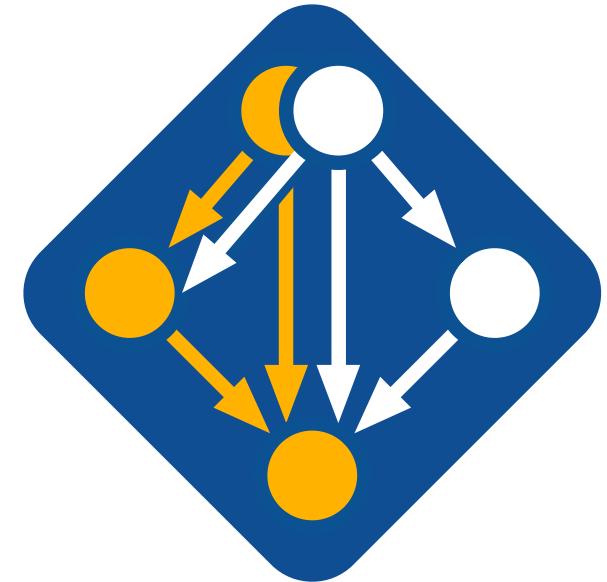
- How to install Spack (works out of the box):

```
$ git clone https://github.com/spack/spack
$ . spack/share/spack/setup-env.sh
```

- How to install a package:

```
$ spack install tau
```

- TAU and its dependencies are installed within the Spack directory.
- Unlike typical package managers, Spack can also install many variants of the same build.
  - Different compilers
  - Different MPI implementations
  - Different build options



Visit [spack.io](http://spack.io)



[github.com/spack/spack](https://github.com/spack/spack)



@spackpm

# Spack provides the *spec* syntax to describe custom configurations

```
$ git clone https://github.com/spack/spack
$ . spack/share/spack/setup-env.sh
$ spack compiler find                                # set up compilers
$ spack external find                               # set up external packages
```

```
$ spack install tau                                unconstrained
$ spack install tau@2.30.1                           @ custom version
$ spack install tau@2.30.1 %gcc@7.3.0              % custom compiler
$ spack install tau@2.30.1 %gcc@7.3.0 +level_zero +/- build option
$ spack install tau@2.30.1 %gcc@7.3.0 +mpi ^mvapich2@2.3~wrapperrpath ^ dependency information
```

- Each expression is a ***spec*** for a particular configuration
  - Each clause adds a constraint to the spec
  - Constraints are optional – specify only what you need.
  - Customize install on the command line!
- Spec syntax is recursive
  - Full control over the combinatorial build space

# E4S v2021-02 GPU Release for x86\_64

```
1: adios2      /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/adios2-2.6.0-nkp24j7enorn3dt7626chuqm3pbkrvfe
2: aml        /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/aml-0.1.0-3mwyb6cf6ervfnruqb5u33v46buuyqth
3: arborx     /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/arborx-0.9-beta-qjzx1kcgplto6pnjpwejoh5xpoik3adr
4: argobots   /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/argobots-1.0-yoafg2slps7kp4dkmb6pzu5z2a37sgs4
5: ascent     /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/ascent-develop-ciwgq6lh6unw3hjsnu47wr7cpqptqgy
6: axom       /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/axom-0.3.3-tzyejxpy3p3ekaev35k2bhpk74cnuh
7: bolt        /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/bolt-1.0-uxku5w5qdfnpa4atgzcabra7wop7lunc
8: caliper    /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/caliper-2.4.0-lfdx3gc6qodg2abbpovib3thdsmsamnn
9: darshan-runtime /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/darshan-runtime-3.2.1-jquggxx2uunyaduo3owhd2snves6mlr
10: dyninst    /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/dyninst-10.2.1-xad3v6rvosm6qfa5fc7d4nn33svtzzf
11: faodel     /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/faodel-1.1906.1-ijilel2vjionmj56mscckw2hpecfsuym
12: flecsi     /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/flecsi-1-c7sevln2ak4pf2jgq6wh3mwictch5l2
13: flit        /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/flit-2.1.0-yvvog7kmax22ei2yyrwdxj3heinmz5am
14: gasnet      /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/gasnet-2020.3.0-ufrq5hym67eq3jsg4jttjjqgo4i6hnq
15: ginkgo      /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/ginkgo-1.2.0-r6lorgchpr5qrcwyqqxtewdhtpi4rmt
16: globalarrays /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/globalarrays-5.7-bow6d32j63j6gusotzjuityznwqv64b
17: gotcha      /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/gotcha-1.0.3-7n7bjnzsfn5w5tnihio3otbaewdhjmu
18: hdf5        /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/hdf5-1.10.6-k74avubedd5knvlc73dr3ib5oyw6bcwn
19: hpctoolkit  /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/hpctoolkit-2020.08.03-wck4g3h3jhfvzxorelxqunbe3xesry
20: hpx         /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/hpx-1.5.0-pynmocntkmuwkowy05jxtycvg34w6kue
21: hypre       /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/hypre-2.19.0-vqo72wn6ei7ruitpg7drkje2rdbdfguo
22: kokkos      /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/kokkos-3.2.0-pqv3uugd6cv3qftyur3rx6dm2gao2tg3
23: kokkos-kernels /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/kokkos-kernels-3.1.00-y4veufyptworlbehxusg4yzh6n7anhp
24: legion      /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/legion-20.03.0-zkbz7h2wuze4dgbwcb04w5fvqltugmog
25: libnrm      /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/libnrm-0.1.0-kp5jb7o4kow25rnggiditwtmdbeebojs
26: libquo      /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/libquo-1.3.1-w45wcw6dqbiajeeaj3ryaeskk7bzx6
27: magma        /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/magma-2.5.3-yksxthffslhjrhzwgcx7smz2tca6ojfn
28: mercury     /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/mercury-1.0.1-ppledtsr3drk2upciytfusuawfxrtjp73q
29: mfem         /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/mfem-4.1.0-kivaike2qintplgufwp5yf2mj3n36ay3
30: mpich        /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/mpich-3.2.1-kgwtpelzbopkrgv24ct6padfbhw7nene
31: mpifileutils /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/mpifileutils-develop-djje5g7ts55g3yic3bms426c2zi7gqsj
32: ninja        /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/ninja-1.10.1-7zbbtuslw25nmqo4ur6abyyf3tchnqv
33: omega-h     /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/omega-h-9.29.0-eln73w7ytpvgqtkmkqyjm4gsabsu2w4p
```

- 67 ECP ST products
- Ubuntu v18.04 x86\_64
- AI/ML package support
  - TensorFlow 2.3.5
  - PyTorch 1.8
  - Horovod
- Support for GPUs
  - Intel oneAPI 2021.1

# E4S v2021-02 GPU Release for x86\_64

```
34: openpmd-api      /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/openpmd-api-0.12.0-4myph6pbjnupgupxdlvbxvqqeqx6atyp
35: openmpi         /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/openmpi-3.1.6-6yqtoym56as6xso2pdgkmn4bcsoyufku
36: papi            /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/papi-6.0.0.1-gorrfrvrik575lldzgq46qmmu63kxl7x
37: papyrus         /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/papyrus-develop-iu3dgpwwyykgv5mpw2dwcr0l4wbwbai
38: parallel-netcdf   /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/parallel-netcdf-1.12.1-tmmkzibn43xr7su76msxxusyzrphdt5
39: pdt              /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/pdt-3.25.1-kvi5wuu5y72fypijti3nxqvdn7zpj6ni
40: petsc            /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/petsc-3.13.4-llg3u4rrt5axrqlim75tt73epewxu4fb
41: plasma           /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/plasma-19.8.1-tji7bojb5ne5hqj2mwn5bqq2tfkm23ke
42: precice          /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/precice-2.1.0-ozdmbat2hlivccha3nklbeahikgynewu
43: pumi             /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/pumi-2.2.2-52czdbxeg7pmjkd55nub5jgxzodcprh
44: py-jupyterhub     /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/py-jupyterhub-1.0.0-tr3wcolaij3kbzb6xm4mbbvakcstwsw3
45: py-libensemble    /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/py-libensemble-0.7.0-mxvqxhiiblqnmlfepbxboyiskqyvbej
46: qthreads          /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/qthreads-1.14-neshsclplh7ttkebm34grztaijqohnxt
47: raja              /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/raja-0.11.0-w25bj2dys6cjqn7isgcjfyvte3tuulev
48: rempi             /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/rempi-1.1.0-sideqdbiik2yseshs3loh4sictbis3t6
49: scr               /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/scr-2.0.0-yh3chyq5gayuk6r4juejjiye6zg3rh3u
50: slate             /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/slate-develop-jnysy2rh5vxhwua5ubtvq4bsfd3py7d5
51: slepc             /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/slepc-3.13.4-q3lalpbqoshiyvjgrnhb2iqiisvnfrp
52: strumpack         /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/strumpack-4.0.0-rlbti5eqc5rjhfisxv2uxevj6m3fn5gg
53: sundials          /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/sundials-5.3.0-3g52gh4a6h4ohucqart5i4m6pi66woj6
54: superlu-dist     /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/superlu-dist-6.3.1-o2hkund66coxn2rrbtlalda2vq35uu7j
55: stc               /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/stc-0.8.3-oxfik7nsmgufoqyy7xilzsRCT7it63ej
56: swig              /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/swig-4.0.1-htxmzjd5sed5yfibw6j7jn5cx6p7g72x
57: sz /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/sz-2.1.9-tcatyiuzh6quctrqd2g3dcli7xa7gvtj
58: tasmanian         /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/tasmanian-7.1-quo3grs5kb2xrvjufpi7vn66cpjfnav
59: tau               /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/tau-2.29-ijw2nbphmlfkt42ubwz7g5a5yru22ikn
60: trilinos          /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/trilinos-13.0.0-6xfnp44g5xm7gpn2en6gkwzfceykdir3x
61: turbine            /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/turbine-1.2.3-q4qjvgxjl3cbuyquo6zrurb4mwfnd6wkp
62: umap              /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/umap-2.0.0-5tob3exzrmwoitudu5pstbb2dms3xnto
63: umpire            /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/umpire-3.0.0-6woo2uvazcucxikc6xad6g3zksu2ygi
64: unifyfs           /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/unifyfs-0.9.0-be7mqbng7kdeewdlgvldhm4jkxnquiil
65: upcxx              /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/upcxx-2020.3.0-pshe62qyvmnrvesqa4pkj6bdq3fxxucf
66: veloc              /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/veloc-1.4-gk3iwfjhmgkwlawp7rmxf2eh37rqpqm2
67: zfp               /opt/spack/opt/spack/linux-ubuntu18.04-x86_64/gcc-7.5.0/zfp-0.5.5-6r6yaco7gga5w4gbuvid3zt2iohrnepj
```

# `spack find` shows what is installed

```
Singularity> spack find
==> 319 installed packages
-- linux-ubuntu18.04-power9le / gcc@7.3.0 --
autoconf@2.69 diffutils@3.7 libiconv@1.16 m4@1.4.18
automake@1.16.2 findutils@4.6.0 libpciaccess@0.16 matio@1.5.17
boost@1.74.0 glm@0.9.7.1 libsigsegv@2.12 metis@5.1.0
bzip2@1.0.8 hdf5@1.10.7 libtool@2.4.6 mpich@3.2.1
cmake@3.18.4 hypre@2.20.0 libxml2@2.9.10 mumps@5.3.3
nurses@6.2 openssl@1.1.1g texinfo@6.5
netcdf-c@4.7.4 parmetis@4.0.3 trilinos@13.0.0
netlib-scalapack@2.1.0 perl@5.26.1 util-macros@1.19.1
omega-h@9.29.0 pkgconf@1.7.3 xz@5.2.5
openblas@0.3.10 suite-sparse@5.7.2 zlib@1.2.11
-- linux-ubuntu18.04-ppc64le / gcc@7.3.0 --
adiak@0.1.1 flit@2.1.0 libpfm4@4.11.0
adios@1.13.1 gasnet@2020.3.0 libpng@1.6.37
adios2@2.6.0 gasnet@2020.3.0 libpthread-stubs@0.4
adlbox@0.9.2 gdbm@1.18.1 libquo@1.3.1
aml@0.1.0 gettext@0.20.2 libsigsegv@2.12
amrex@20.10 ginkgo@1.3.0 libodium@1.0.18
arborx@0.9-beta git@2.28.0 libunistring@0.9.10
argobots@1.0 arpack-ng@3.7.0 libunwind@1.4.0
ascent@develop glm@0.9.7.1 libunwind@1.4.0
autoconf@2.69 globalarrays@5.7 libuuid@1.0.3
automake@1.16.2 gmake@4.2.1 libxml2@2.9.10
axl@0.3.0 gmp@6.1.2 libyogrt@1.24
axom@0.3.3 googletest@1.10.0 libzmq@4.3.2
bash@5.0 gotcha@0.0.2 lmod@8.3
binutils@2.33.1 gotcha@1.0.3 lua@5.3.5
bmi@develop gperf-tools@2.7 lua-luafilesystem@1_7_0_2
bolt@1.0 hdf5@1.8.21 lua-luaposix@33.4.0
boost@1.73.0 hdf5@1.8.21 lwgrp@1.0.3
boost@1.73.0 hdf5@1.10.6 lz4@1.9.2
boost@1.73.0 hdf5@1.10.6 lzo@2.10
boost@1.73.0 boost@1.73.0 hptoolkit@2020.08.03
butterlypack@1.2.0 hpx@1.5.1 magma@2.5.4
bzip2@1.0.8 hwloc@1.11.1 margo@0.4.3
c-blosc@1.17.0 hwloc@2.2.0 matio@1.5.17
caliper@2.4.0 hypre@2.18.2 mbedTLS@2.16.7
cinch@master hypre@2.20.0 mercury@1.0.1
cmake@3.17.3 intel-tbb@2020.3 mercury@1.0.1
conduit@master kokkos@3.2.00 metis@5.1.0
conduit@master kokkos-kernels@3.2.00 mfem@4.1.0
cuda@10.2.89 kvtree@1.0.2 mpark-variant@1.4.0
curl@7.72.0 legion@20.03.0 mpich@3.2.1
darshan-runtime@3.2.1 leveldb@1.22 mpifileutils@develop
da... Snapz Pro X 3.2.1 libarchive@3.4.1 mumps@5.3.3
di... libbsd@0.10.0 ncurses@6.2 py-ipypkgkernel@5.3.4
papyrus@develop py-more-itertools@7.2.0
parallel-netcdf@1.12.1 py-mpi4py@3.0.3
parmetis@4.0.3 py-nbclient@0.5.0
pcre@8.44 py-nbconvert@6.0.1
pcre2@10.35 py-nbformat@5.0.7
pdsh@2.31 py-nest-asyncio@1.4.0
pdt@3.25.1 py-notebook@6.1.4
perl@5.26.1 py-numumpy@1.19.2
petsc@3.13.6 py-oauthlib@3.1.0
petsc@3.14.0 py-pamela@1.0.0
pkgconf@1.7.3 py-pandocfilters@1.4.2
plasma@20.9.20 py-parso@0.6.1
precice@2.1.1 py-petsc4py@3.13.0
pumi@2.2.2 py-pexpect@4.7.0
py-alembic@1.0.7 py-pickleshare@0.7.5
py-argon2-cffi@20.1.0 py-prometheus-client@0.7.1
py-asn1crypto@0.24.0 py-prompt-toolkit@2.0.9
py-async-generator@1.10 py-psutil@5.7.2
py-attrs@19.3.0 py-ptypyprocess@0.6.0
py-babel@2.7.0 py-py@1.8.0
py-backcall@0.1.0 py-pycparser@2.20
py-bleach@3.1.0 py-pyelftools@0.26
py-blinker@1.4 py-pygments@2.6.1
py-certifier@2020.6.20 py-pyjwt@1.7.1
py-certipy@0.1.3 py-pyopenssl@19.0.0
py-cffi@1.14.3 py-pyrsistent@0.15.7
py-chardet@3.0.4 py-pytest-runner@5.1
py-cryptography@2.7 py-python-dateutil@2.8.0
py-cython@0.29.21 py-python-editor@1.0.4
py-decorator@4.4.2 py-python-oauth2@1.1.1
py-defusedxml@0.6.0 py-pytz@2020.1
py-entrypoints@0.3 py-pyzmq@18.1.0
py-idna@2.8 py-requests@2.24.0
py-importlib-metadata@2.0.0 py-send2trash@1.5.0
py-ipypkgkernel@5.3.4 py-setuptools@50.1.0
qthreads@1.14 raja@0.12.1
rankstr@0.0.2 readline@8.0
redset@0.0.3 rempi@1.1.0
scr@2.0.0 shuffle@0.0.3
slate@develop slepc@3.14.0
snappy@1.1.8 sqlite@3.31.1
strumpack@5.0.0 suite-sparse@5.7.2
suite-sparse@5.7.2 sundials@5.4.0
superlu@5.2.1 superlu-dist@6.3.0
superlu-dist@6.3.1 swig@4.0.2
sz@1.4.12.3 sz@2.0.2.0
sz@2.1.10 tar@1.32
tasmanian@7.3 tau@2.29
tcl@8.6.10 texinfo@6.5
turbine@1.2.3 umap@2.1.0
umpire@4.0.1 umpire@4.0.1
unifyfs@0.9.0 unzip@6.0
upcxx@2020.3.0 upcxx@2020.3.0
util-macros@1.19.1
```

All the versions coexist!

- Multiple versions of same package are ok.

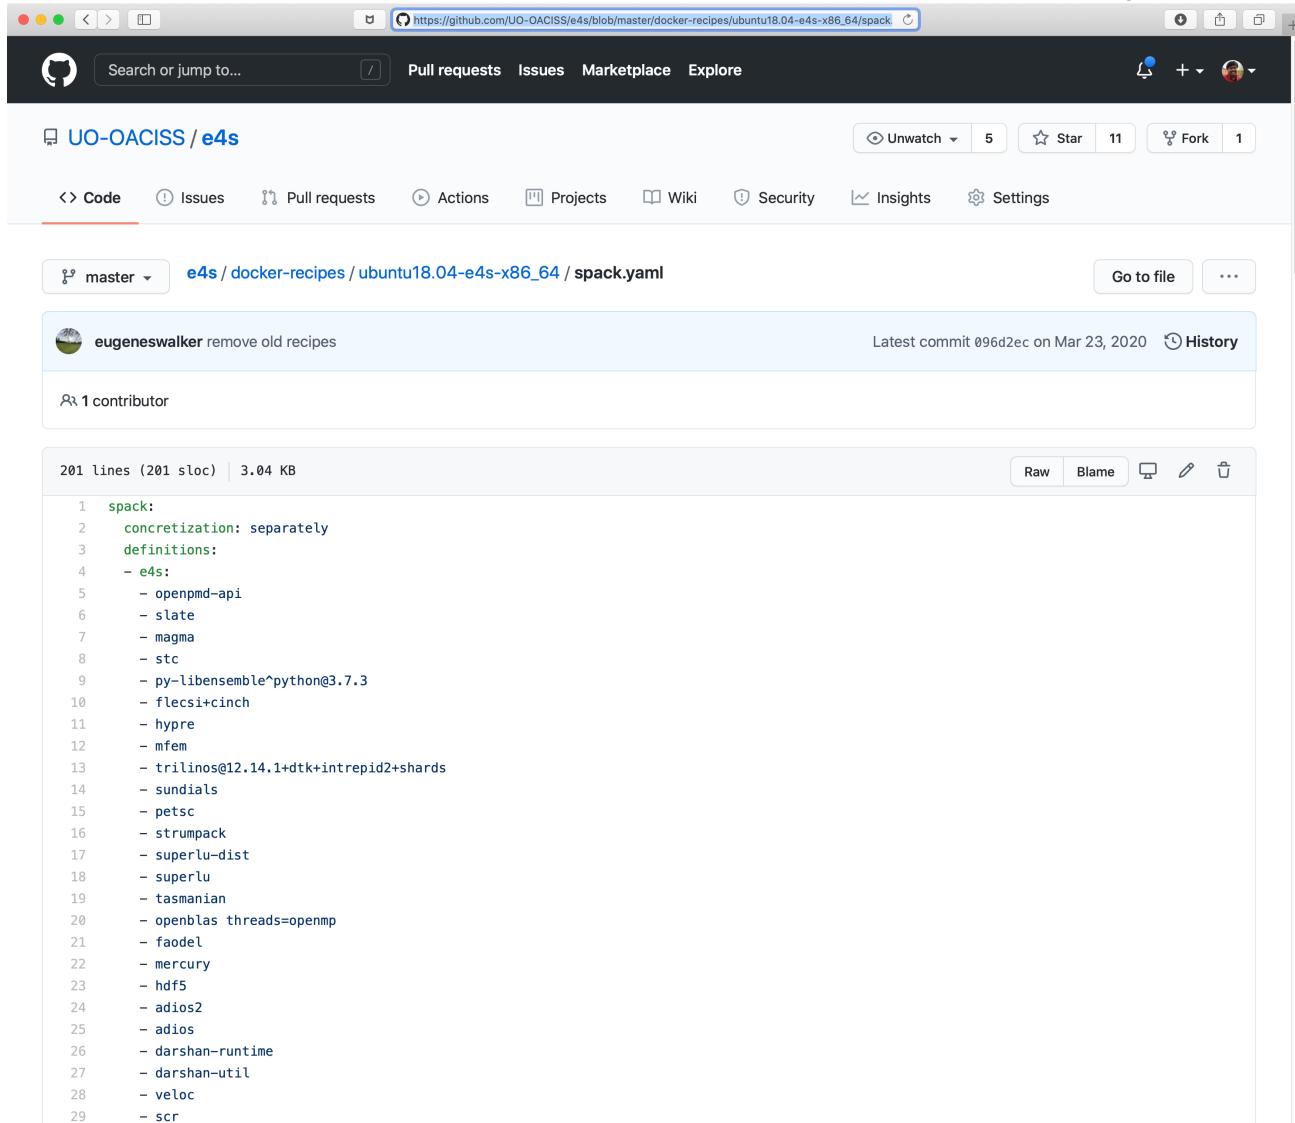
Packages are installed to automatically find correct dependencies.

Binaries work *regardless of user's environment*.

Spack also generates module files.

- Don't have to use them.

# E4S Spack environment spack.yaml



The screenshot shows a GitHub repository page for the `e4s/docker-recipes/ubuntu18.04-e4s-x86_64` branch. The file being viewed is `spack.yaml`. The code listing shows a single block of YAML configuration for a Spack environment named `e4s`, which includes definitions for various packages like `openpmd-api`, `slate`, `magma`, `stc`, `py-libensemble^python@3.7.3`, `flecs+cinch`, `hypre`, `mfem`, `trilinos@12.14.1+dtk+intrepid2+shards`, `sundials`, `petsc`, `strumpack`, `superlu-dist`, `superlu`, `tasmanian`, `openblas threads=openmp`, `faodel`, `mercury`, `hdf5`, `adios2`, `adios`, `darshan-runtime`, `darshan-util`, `veloc`, and `scr`.

```
1 spack:
2   concretization: separately
3   definitions:
4     - e4s:
5       - openpmd-api
6       - slate
7       - magma
8       - stc
9       - py-libensemble^python@3.7.3
10      - flecs+cinch
11      - hypre
12      - mfem
13      - trilinos@12.14.1+dtk+intrepid2+shards
14      - sundials
15      - petsc
16      - strumpack
17      - superlu-dist
18      - superlu
19      - tasmanian
20      - openblas threads=openmp
21      - faodel
22      - mercury
23      - hdf5
24      - adios2
25      - adios
26      - darshan-runtime
27      - darshan-util
28      - veloc
29      - scr
```

- Bare-metal install  
% cat spack.yaml  
% spack -e . install

- Docker build:

Executable File | 2 lines (2 sloc) | 78 Bytes

```
1 #!/bin/bash -x
2 docker build --no-cache -t ecpe4s/ubuntu18.04-e4s-x86_64:1.2 .
```

# E4S Validation Test Suite

- Provides automated build and run tests
- Validate container environments and products
- New LLVM validation test suite for DOE LLVM

The screenshot shows a GitHub repository page for `E4S-Project / testsuite`. The repository has 8 stars and 0 forks. The `validation_tests / magma` branch is selected. A commit by `eugenewalker` is shown, which uses `bash -xe` in `compile/run.sh`. The commit message is "use env variables set by `spack load`". The commit was made 9 hours ago. Below the commit, there is a list of files: `Makefile`, `README.txt`, `clean.sh`, `compile.sh`, `example_f.F90`, `example_sparse.c`, `example_sparse_operator.c`, `example_v1.c`, `example_v2.c`, `run.sh`, and `setup.sh`. The `README.txt` file contains instructions for getting started with MAGMA, including code examples and header inclusion details.

```
use bash -xe in compile/run.sh
...
Makefile      use env variables set by `spack load`        4 months ago
README.txt    Added basic magma test.                      11 months ago
clean.sh      Added basic magma test.                      11 months ago
compile.sh   use bash -xe in compile/run.sh               9 hours ago
example_f.F90 Added basic magma test.                      11 months ago
example_sparse.c Added basic magma test.                  11 months ago
example_sparse_operator.c Added basic magma test.          11 months ago
example_v1.c  Added basic magma test.                      11 months ago
example_v2.c  Added basic magma test.                      11 months ago
run.sh        use bash -xe in compile/run.sh               9 hours ago
setup.sh     Remove some .o files. Don't load special openblas. Don't specify spec... 3 months ago
```

**README.txt**

Getting started with MAGMA.

This is a simple, standalone example to show how to use MAGMA, once it is compiled. More involved examples for individual routines are in the testing directory. The testing code includes some extra utilities that we use for testing, such as `testings.h` and `littest.a`, which are not required to use MAGMA, though you may use them if desired.

-----

C example

See `example_v2.c` for sample code.

Include the MAGMA header:

```
#include "magma_v2.h"
```

(For the legacy MAGMA v1 interface, see `example_v1.c`. It includes `magma.h` instead. By default, `magma.h` includes the legacy cubLAS v1 interface (`cublas.h`). You can include `cublas_v2.h` before `magma.h` if desired.)

# E4S Supports Singularity and Docker container runtimes

Docker images are available on the [E4S Docker Hub](#) and in compressed XZ format on E4S servers.

Recipes for building images from scratch are available on the [E4S GitHub repository](#).

Our recipes make use of Spack packages available as pre-built binaries in the [E4S build cache](#).

 Container Releases

 From source with Spack

[Docker Download](#)

[Singularity x86\\_64 Download](#)

[Singularity ppc64le Download](#)

[CharlieCloud Download](#)

[OVA Download](#)

Spack contains packages for all of the products listed in the E4S 1.0 Full Release category (see above 1.0 Release Notes). General instructions for building software with Spack can be found at the Spack website. For more information, see /usr/local/packages/ecp in the

<https://e4s-project.github.io/download.html>



- wget <http://tau.uoregon.edu/ecp.simg>
- singularity run `ecp.simg`; module load intel
- **Supports Intel OneAPI**
- which dpcpp

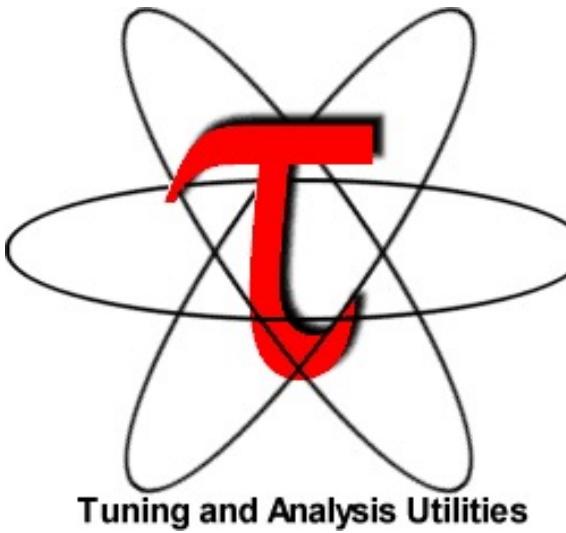
# e4s-cl: A tool to simplify the launch of MPI jobs in E4S containers

- E4S containers support replacement of MPI libraries using MPICH ABI compatibility layer.
- Applications binaries built using E4S can be launched with Singularity using MPI library substitution for efficient inter-node communications.
- e4s-cl is a new tool that simplifies the launch and MPI replacement.
- Usage:

```
. /opt/intel/oneapi/setvars.sh
e4s-cl init --backend singularity --image /home/tutorial/ecp.simg --source /home/tutorial/source.sh
e4s-cl mpirun -np 4 ./a.out
```

<https://github.com/E4S-Project/e4s-cl>

# Download TAU from U. Oregon



<http://tau.uoregon.edu>

for more information

Free download, open source, BSD license

# Performance Research Laboratory, University of Oregon, Eugene



# Support Acknowledgements

- US Department of Energy (DOE)
  - ANL
  - Office of Science contracts, ECP
  - SciDAC, LBL contracts
  - LLNL-LANL-SNL ASC/NNSA contract
  - Battelle, PNNL and ORNL contract
- Department of Defense (DoD)
  - PETTT, HPCMP
- National Science Foundation (NSF)
  - SI2-SSI, Glassbox
- NASA
- CEA, France
- Partners:
  - University of Oregon
  - The Ohio State University
  - ParaTools, Inc.
  - University of Tennessee, Knoxville
  - T.U. Dresden, GWT
  - Jülich Supercomputing Center



# Acknowledgment



“This research was supported by the Exascale Computing Project (17-SC-20-SC), a collaborative effort of two U.S. Department of Energy organizations (Office of Science and the National Nuclear Security Administration) responsible for the planning and preparation of a capable exascale ecosystem, including software, applications, hardware, advanced system engineering, and early testbed platforms, in support of the nation’s exascale computing imperative.”

